

Shape-aware Surface Reconstruction from Sparse Data

Florian Bernard^{1,2,*}, Luis Salamanca², Johan Thunberg², Alexander Tack³, Dennis Jentsch³, Hans Lamecker³, Stefan Zachow³, Frank Hertel¹, Jorge Goncalves², Peter Gemmar^{2,4}

Abstract

The reconstruction of an object’s shape or surface from a set of 3D points is a common topic in materials and life sciences, computationally handled in computer graphics. Such points usually stem from optical or tactile 3D coordinate measuring equipment. Surface reconstruction also appears in medical image analysis, e.g. in anatomy reconstruction from tomographic measurements or the alignment of intra-operative navigation and preoperative planning data. In contrast to mere 3D point clouds, medical imaging yields contextual information on the 3D point data that can be used to adopt prior information on the shape that is to be reconstructed from the measurements. In this work we propose to use a statistical shape model (SSM) as a prior for surface reconstruction. The prior knowledge is represented by a point distribution model (PDM) that is associated with a surface mesh. Using the shape distribution that is modelled by the PDM, we reformulate the problem of surface reconstruction from a probabilistic perspective based on a Gaussian Mixture Model (GMM). In order to do so, the given measurements are interpreted as samples of the GMM. By using mixture components with anisotropic covariances that are oriented according to the surface normals at the PDM points, a surface-based fitting is accomplished. By estimating the parameters of the GMM in a maximum a posteriori manner, the reconstruction of the surface from the given measurements is achieved. Extensive experiments suggest that our proposed approach leads to superior surface reconstructions compared to Iterative Closest Point (ICP) methods.

Keywords: Sparse shape reconstruction, statistical shape model, point distribution model, Gaussian mixture model, expected conditional maximisation.

1. Introduction

The reconstruction of an object’s shape or surface from a set of 3D points is a common topic in materials and life sciences, which is from a computational point of view typically addressed in computer graphics (Berger et al., 2014). Such points usually stem from optical or tactile 3D coordinate measuring equipment, where non-contact measurements are often disturbed by artefacts like reflections, noise, or outliers, making a proper geometric reconstruction even more difficult.

Surface reconstruction also appears in medical image analysis, e.g. in anatomy reconstruction from tomographic measurements via image segmentation. Conceptually speaking, images provide implicit information on the location of anatomical structures via intensity levels. This information is – in many approaches – converted into geometric information via some kind of feature extraction method, either automatically (atlas- or model-based approaches, learning methods, landmark detection, etc.) or

interactively (e.g. via semi-automated or manual labelling of image regions). In both cases, the uncertainty due to noise, artefacts, etc. in the input image data is carried over to the geometric data during the transformation process. In those approaches the segmentation problem is effectively re-phrased as a surface reconstruction problem. Another example from the medical field of computer-assisted surgery is to transfer a pre-operative therapy plan to the operating room (OR) by means of a navigation system (or imaging – however, this can be seen again as a segmentation problem). Here, optical or electromagnetic tracking devices generate 3D positions of a patient’s surface intra-operatively. Matching a therapy plan to those measurements amounts to a surface reconstruction and registration problem.

In contrast to mere 3D point clouds that may represent virtually any object, medical image data yield additional contextual information that can be used to adopt prior information on the anatomical structures to be reconstructed. Heckel et al. (2011) make use of the variational interpolation method (Turk and O’Brien, 1999), which essentially uses a general prior on the surface smoothness. Going one step further, for 3D scan completion Pauly et al. (2005); Gal et al. (2007) have considered templates that are matched to the measurements to achieve the surface reconstruction. However, *these methods are limited since the available measurements are assumed to be sufficiently*

*Corresponding author

¹Centre Hospitalier de Luxembourg, Luxembourg City, Luxembourg

²Luxembourg Centre for Systems Biomedicine, University of Luxembourg, Esch-sur-Alzette, Luxembourg

³Zuse Institute Berlin (ZIB), Germany

⁴Trier University of Applied Sciences, Trier, Germany

dense, as pointed out by Berger et al. (2014).

To tackle this limitation, Bernard et al. (2015b) suggested to use a statistical shape model (SSM) (Heimann and Meinzer, 2009) for surface reconstruction. The motivation is that the class of the anatomical structure of interest is known for clinical routine tasks such as segmentation, registration or intra-operative navigation, which in turn enables to use its shape as a geometric prior. With SSMs being rather strong and specific priors, a central question is how robustly and accurately they can be used for surface reconstruction from a sparse set of surface point measurements.

In this work we present a novel, shape-aware surface reconstruction approach, which will be evaluated with respect to that central question. Through a point distribution model (PDM) (Heimann and Meinzer, 2009) we incorporate a prior into our reconstruction framework that captures the likely shape of an object to be extracted. Using such a PDM, we reformulate the problem of surface reconstruction from a probabilistic perspective. To this end, we embed the prior distribution of the SSM parameters explaining plausible shapes into the objective function.

Our evaluation will consider several different anatomical structures and sparse input point scenarios which may occur in different applications. The goal is to get a deeper understanding on the properties, the general applicability, and the limitations of our approach.

The remainder of this article is organised as follows: In Section 2 related work to our methodology is presented, followed by a concise summary of our key contributions in Section 3. Subsequently, in Sections 4 and 5 we familiarise the reader with the necessary background and we precisely specify the scope of our work. In Section 6 several variants of the shape-aware surface reconstruction method are presented, including a time complexity and convergence analysis. Section 7 contains experiments conducted using the proposed methods. Finally, in Section 8 conclusions are drawn.

2. Related Work

A multitude of methods for general surface reconstruction have been presented in literature so far (see e.g. Raya and Udupa (1990); Bolle and Vemuri (1991); Herman et al. (1992); Hoppe et al. (1992); Edelsbrunner and Mücke (1994); Bajaj et al. (1995); Amenta et al. (1998); Bernardini et al. (1999); Treece et al. (2000); Kazhdan et al. (2006); Schroers et al. (2014)). Many of them are summarized and described in the state-of-the-art (STAR) report by Berger et al. (2014). In the remainder of the present section, only those surface reconstruction methods will be discussed that go beyond simple smoothness assumptions and make use of more explicit shape priors.

For the completion of 2D shapes, Guo et al. (2012, 2013) incorporate templates from a shape database as (geometric) prior into a Bayesian framework. Similarly, a database

of 3D shapes is used by Pauly et al. (2005) for completing 3D surface scans. For improving the flexibility compared to static priors, Gal et al. (2007) use a context-specific database of *local* shape priors, where the input data is matched by (dynamically) combining multiple local shape priors into a global prior. As pointed out by Berger et al. (2014), both approaches described by Pauly et al. (2005) and Gal et al. (2007) have the limitation that the point-clouds are assumed to be sufficiently dense. A unified framework for repairing geometry and texture of meshes has been presented by Park et al. (2006). For filling holes in the surface, a context-based geometry filling procedure is employed, where available local patches of the mesh are used to fill its missing parts. The task of obtaining high-resolution 3D meshes from low-quality input is tackled by Shen et al. (2012). Their approach is to dynamically assemble object templates from a database of object parts. As the name suggests, the 3D shape completion methods discussed so far have the common focus on completing (mostly small) missing parts of meshes obtained from 3D scans.

Blanz et al. (2004) have presented a closed-form solution for SSM-based 3D surface reconstruction from a sparse set of points, which however relies on the assumption that such a set of points is already in correspondence to the model. Albrecht et al. (2013) introduced posterior shape models which have the objective to model the distribution of a whole shape given only partial information. A key assumption is that the corresponding model points of the available partial data are known. In their experiments this issue is either solved manually or using the Iterative Closest Point (ICP) method (Besl and McKay, 1992). Similarly, for shape prediction from sparse observations, Blanc and Székely (2012) use a variant of ICP that evaluates multiple initialisations. The shape completion and animation of people (SCAPE) method has been introduced by Anguelov et al. (2005), where one model for pose deformations and one model for shape variations are learned separately. The main objective of this work is the completion of body shapes based on a small number of known positions for some of the model points. Applied to bone models, Rajamani et al. (2007) fit an SSM to a small number of anatomical landmarks that correspond to some of the model points. Baka et al. (2010) fit an SSM to sparse data points that are in correspondence with the model, which is applied to 2D heart datasets. By producing confidence intervals as output, their method is able to incorporate uncertainties of the input data. Instead of using a trained SSM, Lu et al. (2011) formulate a low-rank matrix recovery problem for restoring missing parts of objects in archaeological studies. Considering that a set of (incomplete) objects of the same class is available, and that correspondences between common parts are known, their approach is based on the assumption that all shapes are approximately linearly correlated. *A shortcoming of the methods discussed so far is that they all assume known correspondences.* However, if the objects do not exhibit a

sufficient amount of distinct features, the identification of exact correspondences is very difficult or even infeasible in practice.

In literature, there are several methods published addressing the mentioned problem of (automatically) detecting the correspondence between sparse data points and a model. Due to its simplicity, the ICP algorithm, where correspondences and transformations are estimated in an alternating manner, is a very popular method for the registration of two shapes represented as point sets. Numerous variants of the originally proposed method have been developed by now, e.g. by Rusinkiewicz and Levoy (2001); Segal et al. (2009); Maier-Hein et al. (2012); Bouaziz et al. (2013); Billings et al. (2015). The surface reconstruction method by Stoll et al. (2006), for instance, deforms a given template in order to fit point-cloud data. For doing so, the user defines initial correspondences between the template and a set of points, which are then refined iteratively in an ICP-like manner. Along the lines of Stoll et al. (2006), for knee surgery, Fleute and Lavallée (1998); Fleute et al. (1999) have presented a methodology to find pose and shape deformation parameters in order to fit an SSM to very sparse data points. Their approach resembles ICP due to the alternating closest point estimation and pose/deformation model parameter updates. Similarly, Chan et al. (2003) reconstruct 3D models of bones in orthopaedic surgery based on a sparse set of points obtained from ultrasound. To this end, shapes are repeatedly instantiated using a PDM, where each shape instance is used as input to ICP in order to (rigidly) fit the sparse points. This work has been extended by Barratt et al. (2008), where a PDM defined on a regular grid instead of the shape surface is used. For sparse shape reconstruction in computer-assisted orthopaedic surgery, a three-stage procedure has been proposed by Zheng et al. (2006). In the first stage, using ICP, the pose of the sparse points is adapted such that they best fit the mean shape. Secondly, the shape deformation parameters are estimated for the given correspondences, which are eventually refined using a further deformation based on thin-plate splines.

Alternatively, in the work of Chang and Zwicker (2009), the registration of articulated shapes in two range scans is based on a reduced deformation model defined on a regular grid. In this approach, the deformations are modelled by a convex combination of rigid transformations, where the weights are spatially varying. The registration is performed by alternately updating the rigid transformations and their weights, where closest point correspondences are recomputed after each step.

Even though the literature on applications and variations of statistical shape models is enormous (for an overview the interested reader is referred to the works by Heimann and Meinzer (2009); Cootes and Taylor (1992); van Ginneken et al. (2002)), the use of SSMs for interactive segmentation is still limited. In the work of van Ginneken et al. (2003), a user directly manipulates points of the PDM, which has the disadvantage that the user

needs to estimate the (unknown) corresponding position of the considered model point in the image. This is in particular difficult if the object does not exhibit distinct features. de Bruijne et al. (2004) present the slice-wise SSM-based segmentation of abdominal aortic aneurysms. For that, the current slice's PDM fitting is initialised from the segmentation result of the previous slice, with the option of manually correcting segmentations on a per-slice basis. Liu and Udupa (2009) present Oriented Active Shape Models, where the semi-automatic live-wire technique is coupled with an Active Shape Model. Other authors present tools for the posterior correction of model-based segmentations, such as Timinger et al. (2003); Schwarz et al. (2008); Tan and Acharya (2014). An interactive procedure based on a set of rules and 2D sketches for completing partial 3D models in structure-from-motion, is presented by van den Hengel et al. (2007).

3. Main Contributions

In this work we present a solution for the reconstruction of surfaces from a sparse set of surface points employing a statistical shape model as a prior. With respect to the solutions already proposed in literature, the novel key contributions of this work can be summarised as follows:

- By extending powerful existing point-set registration procedures (Myronenko et al., 2007; Myronenko and Song, 2010; Horaud et al., 2011), we obtain a shape-aware surface reconstruction method that does *not require explicit correspondences* between the PDM points and the sparse points.
- We recast the matching problem between a sparse set of points and a PDM into a probabilistic framework using Gaussian Mixture Models (GMMs), which we show to be superior to ICP with respect to robustness and accuracy.
- In earlier work we used a spherical (or isotropic) GMM that accounted for a *point-based* matching (Bernard et al., 2015b). We now complement this work by presenting a formulation that is based on anisotropic GMMs that are oriented by the surface normals, which in turn accounts for a *surface-based* fitting.
- A *rigorous derivation* of the surface-based fitting method is presented, leading to an Expected Conditional Maximisation (ECM) algorithm. ECM shares the same convergence properties as the Expectation Maximisation (EM) method while being more general.
- We develop a fast approximation of the ECM-based fitting method. The approximation method has the same computational complexity as the spherical GMM-based method and is empirically less prone to unwanted local optima compared to the ECM-based method.

4. Background

In this section we introduce our notation and we provide a definition of PDMs as they are used throughout the subsequent sections.

4.1. Notation

$\mathbf{1}_p$ and $\mathbf{0}_p$ denote the p -dimensional column vector containing all ones and all zeros, respectively. \mathbf{I}_p denotes the $p \times p$ identity matrix and $\text{diag}(\mathbf{x})$ is the $p \times p$ matrix with the elements of the vector $\mathbf{x} \in \mathbb{R}^p$ on its diagonal. For a matrix \mathbf{A} , by $\mathbf{A}_{i,j}$ we denote the (scalar) element in row i and column j . The colon-notation is used to denote all rows or columns, e.g. $\mathbf{A}_{:,j}$ denotes the column vector formed by the j -th column of \mathbf{A} .

By $\mathbf{p}(x)$ we denote the *probability density function* (p.d.f.), or *probability mass function* in the discrete case, where x can indicate both a random variable and a realisation of it, depending on the context. $\mathcal{N}(x|\mu, \Sigma)$ denotes the Gaussian p.d.f. with mean μ and covariance matrix Σ .

4.2. Point Distribution Models

Statistical shape models based on point distribution models (Heimann and Meinzer, 2009) are an established technique to capture nonlinear shape deformations in \mathbb{R}^3 from training data by using a linear model in the higher-dimensional shape space. Let $\{\mathbf{X}_k : 1 \leq k \leq K\}$ be the set of K training shapes, where each shape is represented by N points (or vertices) in 3 dimensions written as the matrix $\mathbf{X}_k \in \mathbb{R}^{N \times 3}$. In order to process multiple shapes in a meaningful way, it is essential that the rows for all \mathbf{X}_k are corresponding, i.e. the N vertices of all K shapes are in correspondence. The most popular PDM is based on finding an (affine) subspace close to the subspace spanned by the training shapes, which is commonly performed by Principal Components Analysis (PCA).

We first define $\mathbf{X} = [\mathbf{x}_1 \ \dots \ \mathbf{x}_K] \in \mathbb{R}^{3N \times K}$, with $\mathbf{x}_k = \text{vec}(\mathbf{X}_k) \in \mathbb{R}^{3N}$. This allows to compute the modes of shape variability as the eigenvectors of the sample covariance matrix $\mathbf{C} = \frac{1}{1-K}(\mathbf{X} - \bar{\mathbf{x}}\mathbf{1}_K^T)(\mathbf{X} - \bar{\mathbf{x}}\mathbf{1}_K^T)^T$, where $\bar{\mathbf{x}} = \frac{1}{K} \sum_{k=1}^K \mathbf{x}_k$ denotes the mean of all corresponding vertices for all shapes in $3N$ -dimensional shape space. Let $\Phi \in \mathbb{R}^{3N \times M}$ be the matrix of the first M eigenvectors of \mathbf{C} with largest eigenvalues. For a fixed α , the PDM $\mathbf{y}(\alpha) : \mathbb{R}^M \rightarrow \mathbb{R}^{3N}$ is a vector-valued function defined as

$$\mathbf{y}(\alpha) = \bar{\mathbf{x}} + \Phi\alpha, \quad (1)$$

where $\alpha \in \mathbb{R}^M$ is the shape deformation parameter. The deformation of vertex i through α is denoted by

$$y_i(\alpha) = \bar{x}_i + \Phi_i\alpha, \quad (2)$$

where the three rows of vertex i are selected appropriately from $\mathbf{y}(\alpha)$, $\bar{\mathbf{x}}$ and Φ to obtain $y_i(\alpha) \in \mathbb{R}^3$, $\bar{x}_i \in \mathbb{R}^3$ and $\Phi_i \in \mathbb{R}^{3 \times M}$.

A common assumption is that α follows a zero-mean Gaussian distribution, i.e. $\mathbf{p}(\alpha) = \mathcal{N}(\alpha|\mathbf{0}_M, \Lambda)$, where $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_M)$ with λ_m being the m -th eigenvalue of \mathbf{C} . Therefore, thanks to imposing a distribution upon α , we obtain a distribution over shapes (Albrecht et al., 2013).

Usually, in addition to the PDM accounting for shape deformations, a rigid transformation is employed in order to account for the absolute pose of the shape $\mathbf{y}(\alpha)$ with respect to some reference (e.g. the image coordinate system). In the following we assume that the predominant part of the pose has already been normalised. In one of our applications this is motivated by a normalisation of the brain shapes to the ICBM 152 (Fonov et al., 2009) template space (MNI) based on rigid image registration, which allows us to circumvent to explicitly model pose variations. Furthermore, minor pose variations can still be (approximately) captured implicitly in the PDM.

5. Problem Formulation

Given is a PDM that serves as prior for plausible shapes. Also, we assume that (for fixed α) the PDM points are vertices of a (triangular) surface mesh with topology \mathcal{T} . Additionally, we are given the set $\mathcal{P} = \{p_j : 1 \leq j \leq P\}$ of P surface points of the shape that is to be reconstructed, where \mathcal{P} is sparse in the sense that it only contains few measurements p_j of the object's surface. The objective is to find the deformation parameter α such that $\mathbf{y}(\alpha)$ "best agrees" with the points in \mathcal{P} .

5.1. A Generative Model

Motivated by the well-established Coherent Point Drift (CPD) approach for general point-set registration (Myronenko et al., 2007; Myronenko and Song, 2010), we now introduce a probabilistic model based on imposing a distribution over each vertex i by using a GMM.

Given the set \mathcal{P} and a PDM, we consider the following assumptions:

1. For $i = 1, \dots, N$, each PDM vertex position $y_i(\alpha) \in \mathbb{R}^3$ is considered as the mean of a 3D Gaussian distribution with covariance Σ_i .
2. Each point $p_j \in \mathcal{P}$ can be uniquely mapped to one specific vertex index i , its *generating component*, from whose distribution it is drawn.

As such, each point p_j follows the distribution

$$\mathbf{p}(p_j|i, \alpha, \Sigma_i) = \mathcal{N}(p_j|y_i(\alpha), \Sigma_i), \quad (3)$$

where all p_j for $j = 1, \dots, P$ are independent. The corresponding graphical model is depicted in Fig. 1. Note that all of the N Gaussian components are parameterised by α and the covariances $\hat{\Sigma} := \{\Sigma_i\}$. We assume that each component with index i is chosen with equal probability, i.e. $\mathbf{p}(i) = \frac{1}{N}$.




Figure 1: Graphical model of the generating process of \mathcal{P} .

Our objective is to find the parameters α and $\hat{\Sigma}$ that are most likely to have generated the data points \mathcal{P} . Since the generating component i_j of p_j is unknown, the indices of the generating components are treated as *latent* variables. Therefore, to incorporate the uncertainty about the generating component, we consider a GMM for the distribution of p_j , leading to

$$\mathbf{p}(p_j|\alpha, \hat{\Sigma}) = \sum_{i=1}^N \mathbf{p}(i) \mathbf{p}(p_j|i, \alpha, \Sigma_i) \quad (4)$$

$$= \frac{1}{N} \sum_{i=1}^N \mathcal{N}(p_j|y_i(\alpha), \Sigma_i). \quad (5)$$

Using Bayes' theorem, one can derive the probability that the i -th mixture component has generated the point p_j , given also α and $\hat{\Sigma}$, as

$$\mathbf{p}(i|p_j, \alpha, \hat{\Sigma}) = \frac{\mathbf{p}(p_j|i, \alpha, \hat{\Sigma})\mathbf{p}(i)\mathbf{p}(\alpha)\mathbf{p}(\hat{\Sigma})}{\sum_{i'} \mathbf{p}(p_j|i', \alpha, \hat{\Sigma})\mathbf{p}(i')\mathbf{p}(\alpha)\mathbf{p}(\hat{\Sigma})} \quad (6)$$

$$= \frac{\mathbf{p}(p_j|i, \alpha, \Sigma_i)\mathbf{p}(i)}{\sum_{i'} \mathbf{p}(p_j|i', \alpha, \Sigma_{i'})\mathbf{p}(i')} \quad (7)$$

$$= \frac{\mathcal{N}(p_j|y_i(\alpha), \Sigma_i)}{\sum_{i'} \mathcal{N}(p_j|y_{i'}(\alpha), \Sigma_{i'})}. \quad (8)$$

5.2. Optimisation using EM

If the generating component i_j of p_j is unknown, according to eq. (5) all points $p_j \in \mathcal{P}$ are independent and identically distributed (i.i.d.). Thus, the log-likelihood as a function of the model parameters $\theta := (\alpha, \hat{\Sigma})$ reads

$$\begin{aligned} L(\theta) &= \ln \mathbf{p}(\mathcal{P}|\theta) = \ln \prod_{j=1}^P \sum_{i=1}^N \mathbf{p}(i)\mathbf{p}(p_j|i, \theta) \\ &= \sum_{j=1}^P \ln \sum_{i=1}^N \mathbf{p}(i)\mathbf{p}(p_j|i, \theta). \end{aligned} \quad (9)$$

The maximisation of L w.r.t. θ cannot be solved readily due to the sum appearing inside the logarithm. Therefore, a common approach is to employ an iterative method for

the maximisation, where by $\theta^{(n)}$ we denote the estimate of θ at iteration n . Rewriting eq. (9) as

$$L(\theta, \theta^{(n)}) = \sum_{j=1}^P \ln \sum_{i=1}^N \mathbf{p}(i|p_j, \theta^{(n)}) \frac{\mathbf{p}(i)\mathbf{p}(p_j|i, \theta)}{\mathbf{p}(i|p_j, \theta^{(n)})} \quad (10)$$

allows applying Jensen's inequality (Boyd and Vandenberghe, 2009), leading to

$$\begin{aligned} L(\theta, \theta^{(n)}) &\geq \\ \sum_{j=1}^P \sum_{i=1}^N \mathbf{p}(i|p_j, \theta^{(n)}) \ln \frac{\mathbf{p}(i)\mathbf{p}(p_j|i, \theta)}{\mathbf{p}(i|p_j, \theta^{(n)})} &=: Q(\theta, \theta^{(n)}). \end{aligned} \quad (11)$$

As such, the right-hand side of eq. (11), which we denote by $Q(\theta, \theta^{(n)})$, is a lower bound for $L(\theta, \theta^{(n)})$. Maximising this lower bound is the key idea of the well-established EM algorithm (Dempster et al., 1977). In the E-step, the probabilities $\mathbf{p}(i|p_j, \theta^{(n)})$ are evaluated for fixed $\theta^{(n)}$ by using eq. (8). Then, in the M-step, Q in eq. (11) is maximised w.r.t. θ for the fixed $\mathbf{p}(i|p_j, \theta^{(n)})$ computed before. An appealing characteristic of EM is that in each iteration the log-likelihood L is guaranteed to not decrease, which guarantees convergence.

6. Methods

In the following we start by introducing an isotropic GMM, in analogy to the CPD algorithm (Myronenko et al., 2007; Myronenko and Song, 2010). The difference to CPD is that we use a PDM as transformation model. Subsequently, an anisotropic GMM approach is described, where the covariance matrices are oriented according to the surface normals at the PDM points. With that, we move from a purely point-based matching to a more surface-oriented fitting approach. Eventually, we describe a fast approximation of the anisotropic GMM method. By using a simple extension to this approximation, one can ensure that it is an instance of the Generalised Expectation Maximisation (GEM) method and thus the convergence is guaranteed.

6.1. Surface Reconstruction using an Isotropic GMM

Assuming that all N Gaussian components share the same (spherical) covariance $\Sigma_i = \sigma^2 \mathbf{I}_3$, the generative model of eq. (3) reads

$$\mathbf{p}(p_j|i, \alpha, \sigma) = \mathcal{N}(p_j|y_i(\alpha), \sigma^2 \mathbf{I}). \quad (12)$$

With that, the update in the E-step in eq. (8) becomes

$$\mathbf{p}(i|p_j, \alpha, \sigma) = \frac{\exp(-\frac{1}{2\sigma^2}\|p_j - y_i(\alpha)\|^2)}{\sum_{i'} \exp(-\frac{1}{2\sigma^2}\|p_j - y_{i'}(\alpha)\|^2)}. \quad (13)$$

The M-step comprises maximising

$$\begin{aligned} Q(\alpha, \sigma, \alpha^{(n)}, \sigma^{(n)}) &= \text{const} + \\ \sum_{j=1}^P \sum_{i=1}^N \mathbf{p}(i|p_j, \alpha^{(n)}, \sigma^{(n)}) \ln[\mathbf{p}(i)\mathbf{p}(p_j|i, \alpha, \sigma)] &\quad (14) \end{aligned}$$

w.r.t. α and σ . Setting the gradient of Q w.r.t. α to zero reveals that α can be found by solving the linear system $\mathbf{A}\alpha = \mathbf{b}$, with

$$\mathbf{A} = \sum_{i,j} \mathbf{p}(i|p_j, \alpha^{(n)}, \sigma^{(n)}) \Phi_i^T \Phi_i \quad \text{and} \quad (15)$$

$$\mathbf{b} = \sum_{i,j} \mathbf{p}(i|p_j, \alpha^{(n)}, \sigma^{(n)}) \Phi_i^T (p_j - \bar{x}_i). \quad (16)$$

Similarly, σ is computed by

$$\sigma^2 = \frac{1}{3P} \sum_{i,j} \mathbf{p}(i|p_j, \alpha^{(n)}, \sigma^{(n)}) \|p_j - y_i(\alpha)\|^2. \quad (17)$$

Since the maximisation of Q w.r.t. α does not depend on σ , we first compute α and then compute σ with the obtained α , which results in the global maximiser of Q in each M-step. As such, this case is an instance of the ordinary EM algorithm (Dempster et al., 1977). The pseudocode is presented in Algorithm 1.

```

Input:  $\bar{\mathbf{x}}, \Phi, \mathcal{P}$ 
Output:  $\alpha, \sigma^2$ 
Initialise:  $\alpha = \mathbf{0}, \sigma^2 = \frac{1}{3NP} \sum_{i,j} \|p_j - \bar{x}_i\|^2, \mathbf{P} \in \mathbb{R}^{P \times N},$ 
                $\mathbf{y} = \bar{\mathbf{x}} + \Phi\alpha$ 
1 repeat
   // E-step
2   foreach  $j = 1, \dots, P$  do
3      $z = 0$ 
4     foreach  $i = 1, \dots, N$  do
5        $\mathbf{P}_{ji} = \exp(-\frac{1}{2\sigma^2} \|p_j - y_i\|^2)$ 
6        $z = z + \mathbf{P}_{ji}$ 
7      $\mathbf{P}_{j,:} = \frac{1}{z} \mathbf{P}_{j,:}$ 
   // M-step
8    $\alpha = (\sum_i \Phi_i^T \Phi_i \sum_j \mathbf{P}_{ji})^{-1} (\sum_i \Phi_i^T \sum_j \mathbf{P}_{ji} (p_j - \bar{x}_i))$ 
9    $\mathbf{y} = \bar{\mathbf{x}} + \Phi\alpha$ 
10   $\sigma^2 = \frac{1}{3P} \sum_{i,j} \mathbf{P}_{ji} \|p_j - y_i\|^2$ 
11 until convergence

```

Algorithm 1: Pseudocode of the isotropic GMM fitting method.

6.1.1. Maximum A Posteriori (MAP) Solution

We can cast the log-likelihood from eq. (9) into a Bayesian view, leading to

$$L^{\text{posterior}}(\theta) = \ln[\mathbf{p}(\mathcal{P}|\theta)\mathbf{p}(\theta)], \quad (18)$$

where additional knowledge in the form of the prior distribution $\mathbf{p}(\theta)$ is incorporated. By assuming $\mathbf{p}(\alpha) = \mathcal{N}(\alpha|\mathbf{0}_M, \Lambda)$ as described in Section 4.2, and choosing a uniform prior for $\mathbf{p}(\sigma)$, the prior distribution results in $\mathbf{p}(\theta) \propto \mathbf{p}(\alpha)$. When estimating the parameters in the M-step, α is now found by solving the system $\hat{\mathbf{A}}\alpha = \mathbf{b}$ where $\hat{\mathbf{A}} = \mathbf{A} + \sigma^2 \Lambda^{-1}$. The additional term $\sigma^2 \Lambda^{-1}$ corresponds to a Tikhonov regulariser. However, now both updates of α and σ depend on each other, leading to an instance of the ECM algorithm (Meng and Rubin, 1993), which has similar convergence properties as the ordinary EM. To obtain the MAP solution, we have to change the expression in line 8 in Algorithm 1, which uses now the value of σ from the previous iteration to update α . We refer to the MAP solution of the surface reconstruction using an isotropic GMM as ISO.

6.2. Surface Reconstruction using an Anisotropic GMM

The previous approaches use the same (spherical) covariance for each of the N Gaussian components, which can be seen as a purely point-based fitting. However, in a vast amount of medical applications of SSMs, the points of the PDM represent the vertices of a *surface mesh*. This surface mesh is in general only an approximation of a *continuous surface*. Whilst the sparse points \mathcal{P} are assumed to lie on this continuous surface, in general they do not coincide with the PDM vertices. Hence, matching the *surface*, depending on the PDM deformation parameter α , is more appropriate. A didactic 2D example is presented in Fig. 2.

6.2.1. Surface-aligned Covariance Matrices

We now formalise this surface-based fitting method using a GMM with anisotropic covariance matrices that are oriented according to the surface normals. In the GMM, the covariance matrix of each component i , i.e. each vertex of the PDM, is defined as $\Sigma_i(\sigma, \alpha) := \sigma^2 \mathbf{C}_i(\alpha)$. The scalar parameter σ^2 can be seen as a global scaling factor, whereas the matrix $\mathbf{C}_i(\alpha)$ models the anisotropy of the surface structure locally by using a larger variance in the




Figure 2: Anisotropic covariance matrices to achieve a surface-based fitting. The sparse points \mathcal{P} are shown in red, the PDM of a rectangle is shown in green, where the green points define the PDM vertices ($N = 12, M = 2, P = 12$). The orientation of the covariance matrices is shown as white ellipse. The objective is to deform the rectangle PDM such that it fits the red points by adjusting α . The initialisation is shown in (a). Since the red points are sampled between the PDM points (cf. *shape approximation problem* (Hill et al., 1995)), using spherical covariance matrices results in a fit that is even worse than the initialisation (b), whereas using anisotropic covariances results in a more accurate fit (c).

direction of vectors lying in the tangent plane of the PDM surface, compared to the variance along the PDM normal direction (cf. Fig. 2 (c)).

Assuming that the surface topology \mathcal{T} of the underlying shape of the PDM is given in the form of oriented triangles (cf. Section 5), with i_2 and i_3 we denote the index of the “left” and “right” neighbour vertex of i , respectively. With that, the surface normal at vertex i is given by

$$n_i(\alpha) = \frac{(y_{i_2}(\alpha) - y_i(\alpha)) \times (y_{i_3}(\alpha) - y_i(\alpha))}{\|(y_{i_2}(\alpha) - y_i(\alpha)) \times (y_{i_3}(\alpha) - y_i(\alpha))\|}. \quad (19)$$

The covariance matrix of vertex i is defined as

$$\mathbf{C}_i(\alpha) = \left(\frac{1}{\eta} - 1\right)n_i(\alpha)n_i^T(\alpha) + \mathbf{I}_3, \quad (20)$$

where the parameter $\eta \geq 1$ weights the variances of vectors along the normal direction compared to the tangential direction. For the covariance matrix $\Sigma_i(\sigma, \alpha)$, the variance along the normal $n_i(\alpha)$ is given by $\frac{\sigma^2}{\eta}$, and the variance in the direction of any vector in the tangent plane is σ^2 . As such, for $\eta = 1$ one obtains the isotropic GMM described in Section 6.1, and choosing $\eta > 1$ achieves the desired behaviour of modelling a larger variance in the tangent plane. A derivation for eq. (20) can for example be found in the work of Hill et al. (1995). Note that for $\eta > 0$, the matrix $\mathbf{C}_i(\alpha)$ is symmetric and positive definite with the eigenvalues $\{\frac{1}{\eta}, 1, 1\}$. The inverse of \mathbf{C}_i , the *precision matrix*, is given by

$$\mathbf{W}_i(\alpha) := \mathbf{C}_i^{-1}(\alpha) = (\eta - 1)n_i(\alpha)n_i^T(\alpha) + \mathbf{I}_3. \quad (21)$$

6.2.2. MAP Solution

The Q function for the MAP solution now reads

$$\begin{aligned} Q(\alpha, \sigma, \alpha^{(n)}, \sigma^{(n)}) = & \quad (22) \\ & \text{const} - \frac{1}{2}\alpha^T \Lambda^{-1} \alpha - \frac{3P}{2} \ln \sigma^2 \\ & - \frac{1}{2\sigma^2} \sum_{i,j} \mathbf{p}(i|p_j, \alpha^{(n)}, \sigma^{(n)}) \cdot \\ & (p_j - y_i(\alpha))^T \mathbf{W}_i(\alpha) (p_j - y_i(\alpha)). \end{aligned}$$

As already described, the E-step is solved by evaluating eq. (8). Then, the M-step comprises maximising Q in eq. (22) w.r.t. α and σ . Due to the dependence of \mathbf{W}_i on α , finding α that maximises Q does not admit a closed-form solution anymore. Instead, α is now obtained using the BFGS quasi-Newton method (Nocedal and Wright, 2006). The idea is to start with the old value $\alpha^{(n)}$, and then iteratively move along directions that increase Q . Whilst the ordinary Newton method requires the gradient of Q as well as its Hessian, the BFGS quasi-Newton method uses an approximation of the Hessian that is cheap to compute. We now derive the gradient of Q w.r.t. α , i.e.

$$\nabla_{\alpha} Q = \left[\frac{\partial Q}{\partial \alpha_m} \right]_m. \quad (23)$$

For brevity, we write $\partial \cdot$ to denote the partial derivative $\frac{\partial \cdot}{\partial \alpha_m}$ w.r.t. α_m , where the dependence on m is implicit. First, we note that the cross-product $u \times v$ of two vectors $u, v \in \mathbb{R}^3$ can be written as the matrix multiplication $[u]_{\times} v$, where the operator $[\cdot]_{\times} : \mathbb{R}^3 \rightarrow \mathbb{R}^{3 \times 3}$ creates a skew-symmetric matrix from its input vector by

$$[\begin{pmatrix} u_1 \\ u_2 \\ u_3 \end{pmatrix}]_{\times} := \begin{bmatrix} 0 & -u_3 & u_2 \\ u_3 & 0 & -u_1 \\ -u_2 & u_1 & 0 \end{bmatrix}. \quad (24)$$

Introducing

$$b_i(\alpha) := (y_{i_2}(\alpha) - y_i(\alpha)) \times (y_{i_3}(\alpha) - y_i(\alpha)), \quad (25)$$

we can write $n_i(\alpha) = \frac{b_i(\alpha)}{\|b_i(\alpha)\|}$. Now, by representing the cross product in (25) as a matrix product with the notation from (24), and by using the product rule, the partial derivative of $b_i(\alpha)$ is given by

$$\partial b_i(\alpha) = \{ \partial [y_{i_2}(\alpha) - y_i(\alpha)]_{\times} \} (y_{i_3}(\alpha) - y_i(\alpha)) \quad (26)$$

$$\begin{aligned} & + [y_{i_2}(\alpha) - y_i(\alpha)]_{\times} \{ \partial (y_{i_3}(\alpha) - y_i(\alpha)) \} \\ & = [\Phi_{i_2,m} - \Phi_{i,m}]_{\times} (y_{i_3}(\alpha) - y_i(\alpha)) \quad (27) \\ & + [y_{i_2}(\alpha) - y_i(\alpha)]_{\times} (\Phi_{i_3,m} - \Phi_{i,m}). \end{aligned}$$

Moreover,

$$\partial \|b_i(\alpha)\| = \frac{b_i^T(\alpha) \{ \partial b_i(\alpha) \}}{\|b_i(\alpha)\|}. \quad (28)$$

By using the quotient rule, the partial derivative of $n_i(\alpha)$ is given by

$$\partial n_i(\alpha) = \frac{\|b_i(\alpha)\| \{ \partial b_i(\alpha) \} - b_i(\alpha) \{ \partial \|b_i(\alpha)\| \}}{\|b_i(\alpha)\|^2}. \quad (29)$$

Using $\partial n_i(\alpha)$, we can write

$$\partial \mathbf{W}_i(\alpha) = (\eta - 1) (\{ \partial n_i(\alpha) \} n_i^T(\alpha) + n_i(\alpha) \{ \partial n_i^T(\alpha) \}). \quad (30)$$

Now, given the expression for $\partial \mathbf{W}_i(\alpha)$, we can finally compute the partial derivative of Q w.r.t. α_m , which is

$$\begin{aligned} \partial Q(\alpha, \sigma, \alpha^{(n)}, \sigma^{(n)}) = & \quad (31) \\ & - (\Lambda^{-1})_{m,\cdot} \alpha - \frac{1}{2\sigma^2} \sum_{i,j} \mathbf{p}(i|p_j, \alpha^{(n)}, \sigma^{(n)}) \cdot \\ & [(p_j - y_i(\alpha))^T \{ \partial \mathbf{W}_i(\alpha) \} (p_j - y_i(\alpha)) - \\ & 2\Phi_{i,m}^T \mathbf{W}_i(\alpha) (p_j - y_i(\alpha))]. \end{aligned}$$

In order to obtain α , one option is to run the quasi-Newton procedure until convergence, where one obtains an α that (locally) maximises Q . With that, the updates of σ on α depend on each other and the procedure reverts to the ECM algorithm (Meng and Rubin, 1993), as already stated in Section 6.1.1. An alternative is to run only a single quasi-Newton step in each M-step. With that, the obtained α is not a local maximiser of Q ; however, one has

still the guarantee that Q is non-decreasing. As such, this procedure reverts to the GEM algorithm (Dempster et al., 1977).

Finally, the σ -update for fixed α is given by

$$\sigma^2 = \frac{1}{3P} \sum_{i,j} \mathbf{p}(i|p_j, \alpha^{(n)}, \sigma^{(n)}) \cdot (p_j - y_i(\alpha))^T \mathbf{W}_i(\alpha) (p_j - y_i(\alpha)). \quad (32)$$

The pseudocode of the anisotropic GMM fitting procedure is presented in Algorithm 2.

```

Input:  $\bar{\mathbf{x}}, \Phi, \mathcal{P}, \eta, \mathcal{T}$ 
Output:  $\alpha, \sigma^2$ 
Initialise:  $\alpha = \mathbf{0}, \sigma^2 = \frac{1}{3NP} \sum_{i,j} \|p_j - \bar{x}_i\|^2, \mathbf{P} \in \mathbb{R}^{P \times N},$ 
 $\mathbf{y} = \bar{\mathbf{x}} + \Phi \alpha$ 
1 foreach  $i = 1, \dots, N$  do
  // compute  $\mathbf{W}_i$ 
2  $n_i = \frac{(y_{i2} - y_i) \times (y_{i3} - y_i)}{\|(y_{i2} - y_i) \times (y_{i3} - y_i)\|}$ 
3  $\mathbf{W}_i = (\eta - 1)n_i n_i^T + \mathbf{I}_3$ 
4 repeat
  // E-step
5 foreach  $j = 1, \dots, P$  do
6  $t = 0$ 
7 foreach  $i = 1, \dots, N$  do
8  $\mathbf{P}_{ji} = \exp(-\frac{1}{2\sigma^2}(p_j - y_i)^T \mathbf{W}_i(p_j - y_i))$ 
9  $t = t + \mathbf{P}_{ji}$ 
10  $\mathbf{P}_{j,:} = \frac{1}{t} \mathbf{P}_{j,:}$ 
  // M-step
11  $\alpha = \text{quasi-Newton}(Q, \nabla Q, \alpha)$ 
12  $\mathbf{y} = \bar{\mathbf{x}} + \Phi \alpha$ 
13 foreach  $i = 1, \dots, N$  do
  // compute  $\mathbf{W}_i$ 
14  $n_i = \frac{(y_{i2} - y_i) \times (y_{i3} - y_i)}{\|(y_{i2} - y_i) \times (y_{i3} - y_i)\|}$ 
15  $\mathbf{W}_i = (\eta - 1)n_i n_i^T + \mathbf{I}_3$ 
16  $\sigma^2 = \frac{1}{3P} \sum_{i,j} \mathbf{P}_{ji}(p_j - y_i)^T \mathbf{W}_i(p_j - y_i)$ 
17 until convergence

```

Algorithm 2: Pseudocode of the anisotropic GMM fitting method. The notation “quasi-Newton($Q, \nabla Q, \alpha$)” denotes running the quasi-Newton method for maximising Q w.r.t. α , where ∇Q is its gradient and the third argument is the initial value of α . If the GEM approach is used, the quasi-Newton method is run only for a single iteration. Note that the mesh topology \mathcal{T} is used for the normal computations.

In the next section we introduce an approximation of the α -update that is a much faster alternative to the quasi-Newton method.

6.2.3. Fast Approximate Anisotropic GMM

The main idea of our proposed approximation is to use the previous value $\alpha^{(n)}$ instead of α for computing the anisotropic covariance matrices $\mathbf{C}_i(\alpha^{(n)})$ during the α -update in the M-step. Our key assumption is that the PDM is *well-behaved* in the sense that neighbouring vertices vary smoothly during deformation; thus, *locally* the deformation of an individual triangle is nearly a translation. Since surface normals are *invariant* to translations it follows that $\|n_i(\alpha) - n_i(\alpha^{(n)})\|$ is small, which implies that $\|\mathbf{W}_i(\alpha) - \mathbf{W}_i(\alpha^{(n)})\|$ is also small.

The resulting Q function using the proposed approximation is now given by

$$\begin{aligned} \tilde{Q}(\alpha, \sigma, \alpha^{(n)}, \sigma^{(n)}) = & \quad (33) \\ & \text{const} - \frac{1}{2} \alpha^T \Lambda^{-1} \alpha - \frac{3P}{2} \ln \sigma^2 \\ & - \frac{1}{2\sigma^2} \sum_{i,j} \mathbf{p}(i|p_j, \alpha^{(n)}, \sigma^{(n)}) \cdot \\ & (p_j - y_i(\alpha))^T \mathbf{W}_i(\alpha^{(n)}) (p_j - y_i(\alpha)), \end{aligned}$$

where the difference to Q in eq. (22) is that the constant $\mathbf{W}_i(\alpha^{(n)})$ is now used in place of the function $\mathbf{W}_i(\alpha)$. As such, the α -update in the M-step is a quadratic concave problem that can be solved efficiently. The solution for α is found by solving the linear system $\tilde{\mathbf{A}}\alpha = \tilde{\mathbf{b}}$, where $\tilde{\mathbf{A}} \in \mathbb{R}^{M \times M}$ is given by

$$\tilde{\mathbf{A}} = \sigma^2 \Lambda^{-1} + \sum_{i,j} \mathbf{p}(i|p_j, \alpha^{(n)}, \sigma^{(n)}) \Phi_i^T \mathbf{W}_i(\alpha^{(n)}) \Phi_i \quad (34)$$

and $\tilde{\mathbf{b}} \in \mathbb{R}^M$ by

$$\tilde{\mathbf{b}} = \sum_{i,j} \mathbf{p}(i|p_j, \alpha^{(n)}, \sigma^{(n)}) \Phi_i^T \mathbf{W}_i(\alpha^{(n)}) (p_j - \bar{x}_i). \quad (35)$$

The pseudocode for this approximate method is similar to Algorithm 2, except for line 11, where the quasi-Newton method is replaced by solving a linear system for the update of α .

In order to guarantee that the approximate method converges, it is necessary that in the M-step the value of the *exact* Q in eq. (22) is non-decreasing, i.e. the new $\alpha = \tilde{\mathbf{A}}^{-1} \tilde{\mathbf{b}}$ obtained using \tilde{Q} must fulfil

$$Q(\tilde{\mathbf{A}}^{-1} \tilde{\mathbf{b}}, \sigma^{(n)}, \alpha^{(n)}, \sigma^{(n)}) \geq Q(\alpha^{(n)}, \sigma^{(n)}, \alpha^{(n)}, \sigma^{(n)}). \quad (36)$$

For $\eta = 1$ the methods reverts to the isotropic method and condition (36) vacuously holds. However, for $\eta > 1$ this is not true in general. A simple way to ensure that Q is non-decreasing is to evaluate the condition in eq. (36) in each iteration, and in the case of a violation revert to one of the quasi-Newton methods for the α -update. Specifically, we consider a single quasi-Newton step for updating α . We denote the approximate method without this convergence check as *ANISO*, and the method with the convergence check and the quasi-Newton step as fallback as *ANISOc*.

6.3. Performance Analysis

Table 1 summarises the computational complexity of the individual subroutines involved in the presented methods.

In Fig. 3 we plot the mean of the normalised value of Q as a function of the processing time for all four anisotropic fitting methods. For each random run we sample a shape instance by drawing α (cf. Section 4.2), select P points randomly from the mesh surface (cf. Section 7), and run the four methods. The obtained Q for the four methods

Table 1: Computational complexity table. The complexity of the α -update for one iteration of the BFGS quasi-Newton methods is $\mathcal{O}(M^2)$ plus the complexity of the evaluation of Q and ∇Q (Nocedal and Wright, 2006) (we use n to denote the number of iterations of the quasi-Newton method). The complexity of the α -update of the remaining methods comprises the computation of $\mathbf{A}/\tilde{\mathbf{A}}$ and $\mathbf{b}/\tilde{\mathbf{b}}$, as well as solving a linear system of equations of size $M \times M$, for which we present the complexity $\mathcal{O}(M^3)$ due to the matrix inversion involved. Note that in * we present the complexity for general Λ , for diagonal Λ the quadratic time complexity in M reduces to linear complexity.

		ECM	GEM	anisotropic ANISOc	ANISO	ISO
update \mathbf{y}				$\mathcal{O}(MN)$		
compute $\{\mathbf{W}_i\}$				$\mathcal{O}(N)$		-
E-step				$\mathcal{O}(NP)$		
α -update	evaluate Q , eq. (22)		$\mathcal{O}(MN + M^2 + NP)^*$			-
	evaluate ∇Q , eq. (26)		$\mathcal{O}(M^2 + MNP)^*$			-
	construct $\mathbf{A}/\tilde{\mathbf{A}}$		-		$\mathcal{O}(M^2N + NP)$	
	construct $\mathbf{b}/\tilde{\mathbf{b}}$		-		$\mathcal{O}(MN + NP)$	
	total α -update	$\mathcal{O}(n(M^2 + MNP))$	$\mathcal{O}(M^2 + MNP)$	$\mathcal{O}(M^3 + MNP + M^2N)$	$\mathcal{O}(M^3 + M^2N + NP)$	
σ -update				$\mathcal{O}(NP)$		
total (per outer iteration)		$\mathcal{O}(n(M^2 + MNP))$	$\mathcal{O}(M^2 + MNP)$	$\mathcal{O}(M^3 + MNP + M^2N)$	$\mathcal{O}(M^3 + M^2N + NP)$	

are then normalised such that in each run the smallest Q corresponds to 0 and the largest Q corresponds to 1 (normalisation w.r.t to all four methods simultaneously). We have found that the single-step quasi-Newton method (GEM) is faster compared to the full quasi-Newton procedure (ECM). Moreover, compared to both quasi-Newton methods, the approximate methods are much faster. Since the ANISOc method makes use of elements both of the GEM and the ANISO method, the total time complexity of the ANISOc method is the combined time complexity for GEM and ANISO (cf. Table 1). Nevertheless, in our simulations the ANISOc method comes close to the ANISO method in terms of convergence speed. This is because in the early stages of the iterative procedure the ANISOc method in most cases satisfies condition (36) by using approximate M-steps. A violation of condition (36) happens more frequently in the later stages of the iterative procedure, where additionally the slower quasi-Newton M-step update has to be performed.

Surprisingly, for $\eta = 64$, the ANISO method is on average *better* than all three other methods. This can be explained by the fact that with increasing η the function Q becomes “more non-concave”, which introduces more local optima. In the ANISO method we find α by maximising the *concave* function \tilde{Q} instead of the *highly non-concave* function Q . Hence, our simulations suggest that it is advantageous to find a *global* maximum of an approximate Q , i.e. \tilde{Q} , instead of finding a *local* maximum of the exact Q .

In Fig. 4 we illustrate the behaviour of Q and compare it with \tilde{Q} for various choices of η . Note that for visualisation purposes we have chosen $M = 2$, whereas in higher-dimensional cases the effect of an increasing η on the non-concavity can be expected to be more severe.

Since the results shown in Fig. 3 suggest that the ANISO method is faster whilst being at least as good as the




Figure 3: Proportion of best value of Q versus processing time averaged over 100 random runs for the four anisotropic methods for two choices of η . Results were obtained using the brain shapes dataset ($N = 1792, M = 16, P = 30$, cf. Section 7.2).

other methods, for further experiments we chose to use the ANISO method as representative for the anisotropic fitting methods, which allowed us to evaluate more extensive configurations due to its computational efficiency.

7. Experiments

In this section we evaluate the proposed fitting procedures on two datasets. For the generation of the set of sparse points \mathcal{P} , we sample sparse points randomly on the shape surfaces. For sampling the random points on the shape *surface*, we first select a triangle from the surface mesh with a probability proportional to its area. Then,

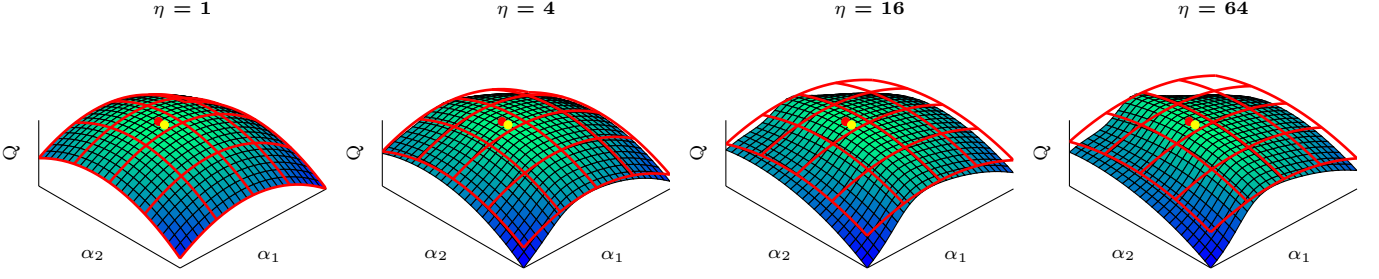


Figure 4: Illustration of the behaviour of Q for various η . The height and the colour of the surface both show the value of Q , eq. (22), depending on α_1 and α_2 . The red grid shows its concave approximation \tilde{Q} as presented in eq. (33). The red dot denotes the value of Q at $\alpha^{(n)}$; at this position $Q = \tilde{Q}$. The yellow dot indicates the maximum of Q . For the trivial case of $\eta = 1$ it can be seen that $Q = \tilde{Q}$ everywhere, whereas an increasing η leads to a larger discrepancy between Q and \tilde{Q} as well as to an “increased non-concavity” of Q .

we uniformly sample a point lying within the triangle according to the procedure presented by Osada et al. (2002). Moreover, we evaluate noisy versions of these points by adding Gaussian noise with standard deviation $\tilde{\sigma}$ to each point individually.

In addition to the probabilistic fitting methods presented in this paper, we also use a simple (regularised) ICP method as baseline for fitting the PDM to the sparse points, which is outlined in Algorithm 3. The anisotropic method requires to set the parameter η accounting for the amount of anisotropy. We evaluate various choices of $\eta = 2, 4, 8, 16$ and report results for the value of η that leads to the highest average Dice Similarity Coefficient (cf. Section 7.1).

We consider leave-all-in (LAI) and leave-one-out (LOO) experiments. The LAI experiments measure the performance of our method given a perfect model, whereas the LOO experiments evaluate the generalisation ability to unseen data.

Input: $\bar{\mathbf{x}}, \Phi, \mathcal{P}, \Lambda$
Output: α
Initialise: $\alpha = \mathbf{0}$, $\mathbf{p} = \text{vec}([p_1, \dots, p_P]) \in \mathbb{R}^{3P}$

```

1 repeat
2    $\mathbf{y} = \bar{\mathbf{x}} + \Phi\alpha$ 
  // Find nearest neighbours
3    $\mathcal{N} = \text{findNearestNeighbourIndices}(\mathbf{y}, \mathcal{P})$ 
  // Solve linear system for  $\alpha$ 
4    $\mathbf{A} = \Phi_{\mathcal{N},:}$ 
5    $\mathbf{b} = \mathbf{p} - \bar{\mathbf{x}}_{\mathcal{N}}$ 
6    $\alpha = (\mathbf{A}^T \mathbf{A} + \Lambda^{-1})^{-1} \mathbf{A}^T \mathbf{b}$  // Tikhonov regularisation
7 until convergence
```

Algorithm 3: Pseudocode of the ICP baseline method. The notation $\Phi_{\mathcal{N},:}$ and $\bar{\mathbf{x}}_{\mathcal{N}}$ means selecting the appropriate rows from Φ and $\bar{\mathbf{x}}$ according to the indices of the nearest neighbours \mathcal{N} .

7.1. Evaluation Metrics

Let \mathcal{X} be the set of points representing the vertices of a high-resolution mesh of the ground truth shape, and let \mathcal{Y}

be the set of the points of the PDM after fitting. We use

$$\mathcal{N}_x = \arg \min_{y \in \mathcal{Y}} \|x - y\|_2 \quad \text{and} \quad (37)$$

$$\mathcal{N}_y = \arg \min_{x \in \mathcal{X}} \|x - y\|_2 \quad (38)$$

to denote the nearest neighbour of x in \mathcal{Y} and of y in \mathcal{X} , respectively.

The (symmetric) Hausdorff distance S_{\max} between \mathcal{X} and \mathcal{Y} is given by

$$S_{\max}(\mathcal{X}, \mathcal{Y}) = \max \left\{ \max_{x \in \mathcal{X}} \|x - \mathcal{N}_x\|_2, \max_{y \in \mathcal{Y}} \|y - \mathcal{N}_y\|_2 \right\}. \quad (39)$$

The (symmetric) average distance S_{avg} is given by

$$S_{\text{avg}}(\mathcal{X}, \mathcal{Y}) = \frac{1}{\|\mathcal{X}\| + \|\mathcal{Y}\|} \left(\sum_{x \in \mathcal{X}} \|x - \mathcal{N}_x\|_2 + \sum_{y \in \mathcal{Y}} \|y - \mathcal{N}_y\|_2 \right). \quad (40)$$

Additionally to both surface-based metrics, we use the following metrics that are based on the correspondence between the points of a pair of shapes $\mathbf{X} \in \mathbb{R}^{N \times 3}$ and $\mathbf{Y} \in \mathbb{R}^{N \times 3}$. The maximum-per-vertex distance D_{\max} is given by

$$D_{\max}(\mathbf{X}, \mathbf{Y}) = \max_{i=1, \dots, N} \|\mathbf{X}_{i,:} - \mathbf{Y}_{i,:}\|_2, \quad (41)$$

and the average-per-vertex distance D_{avg} is given by

$$D_{\text{avg}}(\mathbf{X}, \mathbf{Y}) = \frac{1}{N} \sum_{i=1, \dots, N} \|\mathbf{X}_{i,:} - \mathbf{Y}_{i,:}\|_2. \quad (42)$$

Moreover, we use the Dice Similarity Coefficient (DSC) as volumetric metric, which is defined as

$$\text{DSC}(\mathcal{V}_x, \mathcal{V}_y) = \frac{2|\mathcal{V}_x \cap \mathcal{V}_y|}{|\mathcal{V}_x| + |\mathcal{V}_y|} \quad (43)$$

for the volumetric segmentations \mathcal{V}_x and \mathcal{V}_y .

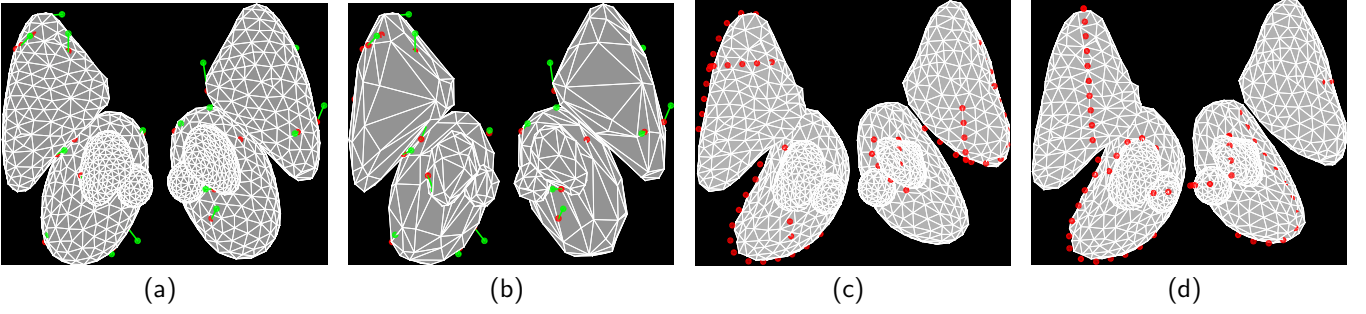


Figure 5: Brain shapes dataset. (a) Mean shape with $N = 1792$ vertices. (b) Downsampled mean shape with $\approx 0.2N$ vertices. In (a,b) $P = 36$ sparse points have been randomly drawn from the original surface according to the procedure described in Section 7. The noise-free points are shown in red, whereas their noisy versions with $\tilde{\sigma} = 2$ are shown in green. (c,d) Shape instances from the training set with contour points generated according to Algorithm 4 in the Appendix are shown, where $n_{\text{obj}} = 4$, $n_{\text{cont}} = 2$ in (c) and $n_{\text{obj}} = 8$, $n_{\text{cont}} = 1$ in (d).

7.2. Brain Shapes

For deep brain stimulation (DBS) surgery the availability of patient-specific 3D models of relevant brain structures would be highly beneficial both for surgery planning and for intra-operative navigation (Bernard et al., 2012). However, a manual segmentation is too costly for the daily clinical routine use and the clinical acceptance of fully computer-generated 3D models for such highly individualised surgeries stems a problem for automated segmentation methods. One interesting application of our presented shape-aware surface reconstruction method is interactive segmentation. This could be implemented by alternating between the user annotating brain structure boundaries and running our fitting method in order to reconstruct a surface from the user-input.

We consider a multi-object PDM that captures the inter-relation between the brain structures *Substantia Nigra & Subthalamic Nucleus* (SN+STN, as compound object), *Nucleus Ruber* (NR), *Thalamus* (Th) and *Putamen & Globus Pallidus* (Put+GP, as compound object), where all structures are considered bilaterally. The mean of the PDM is shown in Fig. 5 (a). The motivation for using a *multi-object* PDM is to be able to reconstruct surfaces of objects, even if for some of the structures no annotation has been provided, e.g. because the image contrast is too low for identifying the boundaries of these particular structures with high confidence.

In Fig. 5 (a), random sparse points that are generated according to the procedure described in Section 7 are shown. In addition, we generate a set \mathcal{P} that comprises (partial) contours, as shown in Fig. 5 (c,d). The idea is to generate these partial contours in a manner as they could be drawn by the user with an interactive pen display. The procedural generation of these points is simple in nature but lengthy, as outlined in Algorithm 4 in the Appendix A. To generate the set \mathcal{P} of points that lie on the partial contours, n_{obj} out of the total number of eight brain structures (four in each hemisphere) are selected. For each of the n_{obj} selected structures, n_{cont} partial contours are generated. A partial contour is generated by using a sub-

contour that has half the length of the full object boundary for a particular slice and orientation. Moreover, the selection of the n_{obj} is not arbitrary; we enforce that at least $n_{\mathcal{L}}$ of them must be chosen from the set of obligatory labels \mathcal{L}' . The latter is done in order to ensure that the resulting problem is well-posed, i.e. to avoid that only structures that are very close together are selected. We evaluated two settings, $(n_{\text{obj}} = 4, n_{\text{cont}} = 2, n_{\mathcal{L}} = 2, \mathcal{L}')$ with $\mathcal{L}' = \{\text{Put+GP}_L, \text{Put+GP}_R, \text{Th}_L, \text{Th}_R\}$ and $(n_{\text{obj}} = 8, n_{\text{cont}} = 1)$; note that for the latter a single (partial) contour is generated for all of the eight brain structures, thus, the specification of \mathcal{L}' is unnecessary. The particular choice of the set of obligatory labels \mathcal{L}' is based on the observation that the structures in \mathcal{L}' encompass the entire multi-object complex. Note that when considering partial contours, for each $p_j \in \mathcal{P}$ we assume that it is known to which of the eight brain structures it belongs, which is used to constrain the E-step in our fitting methods (and the nearest-neighbour routine for the ICP method).

Qualitatively, we found that this procedure generates contour segments that look plausible to have been drawn by a user. In order to evaluate the robustness of our presented method with respect to noisy input, we also created noisy contours. For that, the entire boundary \mathcal{B} is translated by a random vector that has a zero-mean Gaussian distribution with covariance $\tilde{\sigma}\mathbf{I}_2$ (line 18 in Algorithm 4 in the Appendix).

7.2.1. Pose Normalisation

The PDM is learnt from multi-label segmentations that are all represented in a common coordinate system, the MNI template space in our case (more details on the manual annotation and the establishment of correspondences can be found in our previous work (Bernard et al., 2014, 2016)). The alignment of the patient images into the MNI template space is conducted using the rigid image registration method FLIRT (Jenkinson and Smith, 2001). Hence, thanks to this alignment, the orientation and position are already approximately normalised. Nevertheless, minor pose variations may still be present between the individual

images. Whilst the translational part of these pose differences can be captured loss-free by the PDM, the rotational part can only be captured approximately. However, due to the normalisation to the MNI space it can be expected that the rotational part is small and as such the first-order approximation that is captured by the PDM is of sufficient accuracy. Consequently, for a new patient image that is to be segmented, a normalisation to the MNI template space is sufficient.

7.2.2. Results

For each of the $K = 17$ training shapes we sample 20 instances of sparse points \mathcal{P} , leading to $17 \cdot 20 = 340$ runs per method. For the LAI simulations we used a PCA-based PDM with $M = 16$ modes. For the LOO experiments we used kernel PCA (kPCA) with $M = 96$ modes to learn the PDM, which is able to improve the generalisation ability of the PDM from our small training dataset comprising $K = 17$ shapes (Bernard et al., 2015a). First, we compute all metrics for each of the eight brain structure objects individually, which are then summarised by computing the average for DSC , S_{avg} and D_{avg} ; and by computing the maximum for S_{max} and D_{max} .

A summary of the results comparing the ICP method and the ANISO method are shown in Figs. 6 and 7. In each row a different metric is shown. The left column depicts the mean and the right column the standard deviation. Each subplot shows the value of the metric as a function of the number of sparse points P . The solid lines indicate the noise-free case and the dotted lines the noisy case with $\tilde{\sigma} = 1$.

It can be seen that the ANISO method outperforms the ICP method in all cases, where for both the ICP and the ANISO method an increasing number of measurements P improves the results. However, in the LAI case as shown in Fig. 6, the gap between ANISO and ICP is larger in the smaller ranges of P , whereas once more points are available ($P = 90$), the gap between the ICP method and the ANISO method becomes smaller. In the LOO case, as shown in Fig. 7, this does not hold anymore.

In both the LAI and LOO case the standard deviation of ICP is much larger compared to our method. Moreover, using the ANISO method the standard deviation can be reduced by increasing P , which does not hold for the ICP method in the LOO setting.

More detailed results that back up these findings are presented in Figs. B.11-B.20, which can be found in Appendix B.




Figure 6: Summary of brain shape LAI results.

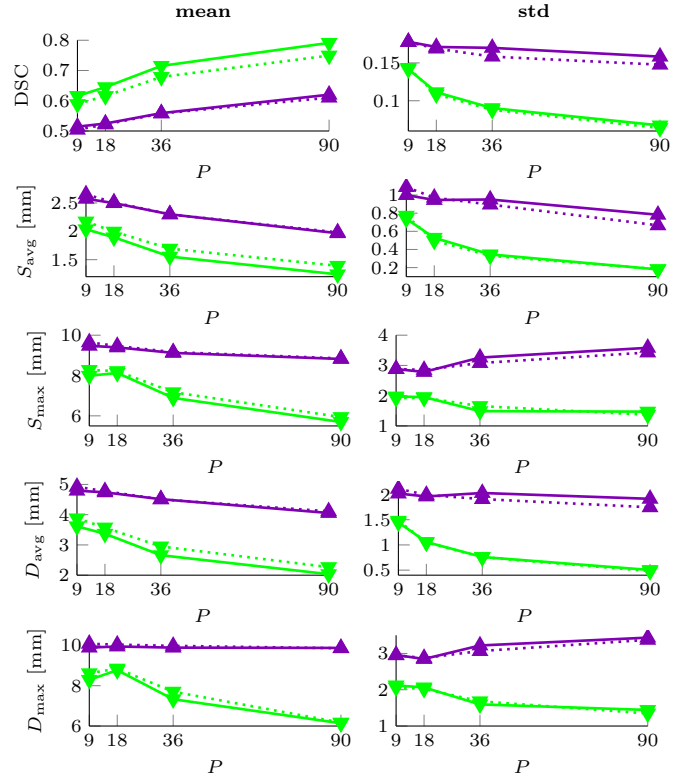


Figure 7: Summary of brain shape LOO results.

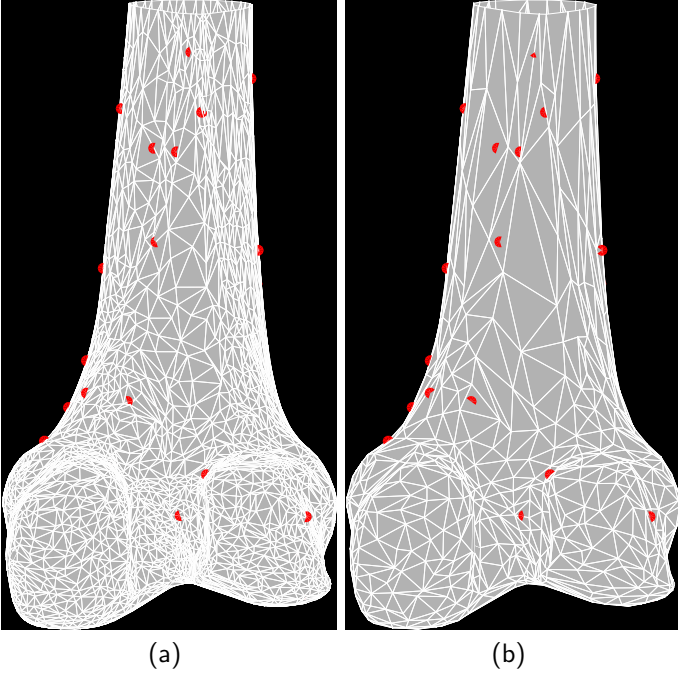


Figure 8: Femur dataset. (a) Mean shape with $N = 3793$ vertices. (b) Downsampled mean shape with $\approx 0.2N$ vertices. $P = 36$ sparse points have been randomly drawn from the original surface according to the procedure described in Section 7.

7.3. Femur

Knee malalignment is assumed to be associated with knee osteoarthritis (OA), where both *Varus* and *Valgus* malalignment are a potential risk factor. In late stages of OA a total knee replacement is often the ultima ratio. In total knee prosthesis malalignment has a direct impact on the load situation. Knowledge of the accurate bone shape aids implant design as well as patient specific planning and could also predict OA progression.

However, the manual segmentation of knee bones is tedious, time consuming and highly subjective. Thus, in recent years many semi-automatic and fully automated segmentation methods have been proposed for CT and MRI data of the human knee (Heimann and Meinzer, 2009; Arovitola and Gallo, 2016). In this section we illustrate the potential benefit of our method for knee bone segmentation for a femur SSM with an analogous application in tibia and patella segmentation. Given points that lie on the femur’s surface, the proposed shape-aware surface reconstruction method can be used for the task of SSM-based femur segmentation. These points could be manually defined or automatically detected by algorithms like SIFT (Lowe, 1999), SURF (Bay et al., 2006) or machine learning techniques in general (Yang et al., 2015; Xue et al., 2015).

7.3.1. Pose Normalisation

For the femur dataset we assumed that the pose has already been normalised and we directly worked in the space

of the SSM. In practice, this can for example be tackled in a similar manner as by Seim et al. (2010), who proposed an automated SSM-based knee bone segmentation, where initially the model is positioned inside the three dimensional CT or MR image via Generalised Hough Transform (Ballard, 1981).

7.3.2. Results

We present experiments for a PDM of the femur with $N = 3793$ points (cf. Fig. 8 (a)). Additionally, we evaluated the ANISO method using a downsampled PDM, denoted ANISO-ds, where a subset of only $\approx 0.2N$ of the original PDM vertices is used (cf. Fig. 8 (b)). Random sparse points are generated according to the procedure described in Section 7, where for each training shape 10 instances of sparse points \mathcal{P} are sampled. In Fig. 8 (a,b) such a random instance of \mathcal{P} is shown.

LAI and LOO experiments have been carried out for the femur model evaluating the same metrics as before. Our PCA-based PDM has $K = 60$ training shapes with $M = 59$ modes for the LAI experiments, and has $K = 59$ training shapes with $M = 58$ modes for the LOO experiments.

A summary of the results comparing the ICP method, the ANISO method and the ANISO-ds method are shown in Figs. 9 and 10. In each row a different metric is shown. The left column depicts the mean and the right column the standard deviation. Each subplot shows the value of the metric as a function of the number of sparse points P .

It can be seen that if only $P = 9$ points are available, the ICP method and the ANISO-ds method produce comparable results. Once more points become available, both anisotropic methods outperform the ICP method. Surprisingly, the ANISO-ds method, which uses a downsampled PDM, outperforms the ANISO method for $P = 9$. We assume that this is because the original PDM with $N = 3793$ vertices contains fine local details that lead to an overfitting when reconstructing the surface from only $P = 9$ points. In contrast, the downsampled PDM contains less details that may impede the surface reconstruction.

Additional plots are shown in Figs. C.21, C.22 in the Appendix C.




Figure 9: Summary of femur LAI results.

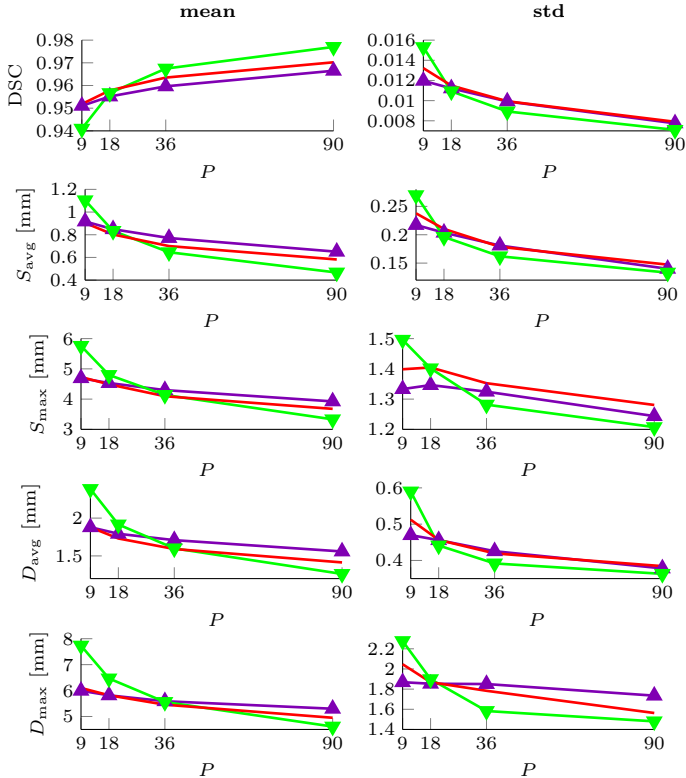


Figure 10: Summary of femur LOO results.

8. Conclusion

In this paper we have presented a methodology for the shape-aware surface reconstruction from sparse surface points. For that, the likely shape of the object that is to be reconstructed is captured by a PDM associated with a surface mesh. By interpreting the available sparse surface points as samples of a GMM, the surface reconstruction task is cast as maximisation of the posterior likelihood, which we tackle by variants of the EM algorithm. In order to achieve a surface-based fitting, we use a GMM with anisotropic covariance matrices, which are oriented by the surface normals at the PDM points. However, this results in a non-concave optimisation problem that is to be solved in each M-step. We deal with this by maximising a concave approximation that considers the surface normals of the PDM computed from the previous value of the shape deformation parameter. Based on the assumption that neighbour PDM vertices vary smoothly and the fact that surface normals are invariant to translations, we have presented an intuitive explanation why this approximation makes sense. Empirically we demonstrated that finding a global maximum of this approximation leads to better results compared to finding a local optimum during the exact (non-concave) M-step. Moreover, our proposed concave approximation results in an algorithm that has the same time complexity as the isotropic fitting procedure.

A limitation of the proposed method is that we do not explicitly consider a rigid transformation in order to model the pose of the object that is to be reconstructed. Whilst in principal one can formulate an analogous problem that considers the pose, the resulting problem is much more difficult to solve. This is because a simultaneous maximisation with respect to the rigid transformation and the shape deformation parameter must be performed, which is usually done iteratively, as for example during Active Shape Model search (Cootes and Taylor, 1992). With that, the resulting surface reconstruction procedure would be much more sensitive to unwanted local optima as well as much slower. Instead, our surface reconstruction method deals with shape deformations only and the normalisation of the pose must be solved a-priori in an application-dependent manner. In our example of the multi-object brain shape reconstruction we dealt with this issue using rigid image registration in order to align the data into a common coordinate system.

Experimentally we have demonstrated that our method leads to superior fitting accuracy compared to ICP for both the brain dataset and the femur dataset. Moreover, the variance of the results obtained using our method is drastically smaller compared to ICP for the brain dataset.

Acknowledgement

The authors gratefully acknowledge the financial support by the Fonds National de la Recherche, Luxembourg (6538106, 8864515, 9169303).

References

References

- Albrecht, T., Lüthi, M., Gerig, T., Vetter, T., 2013. Posterior shape models. *Medical Image Analysis* 17, 959–973.
- Amenta, N., Bern, M., Kamvysselis, M., 1998. A new Voronoi-based surface reconstruction algorithm, in: *SIGGRAPH*.
- Anguelov, D., Srinivasan, P., Koller, D., Thrun, S., Rodgers, J., Davis, J., 2005. SCAPE: shape completion and animation of people, in: *SIGGRAPH*, ACM. pp. 408–416.
- Aprovitola, A., Gallo, L., 2016. Knee bone segmentation from mri: A classification and literature review. *Biocybernetics and Biomedical Engineering* (in press).
- Bajaj, C.L., Bernardini, F., Xu, G., 1995. Automatic reconstruction of surfaces and scalar fields from 3D scans, in: *SIGGRAPH*.
- Baka, N., de Bruijne, M., Reiber, J., 2010. Confidence of model based shape reconstruction from sparse data, in: *Biomedical Imaging: From Nano to Macro*.
- Ballard, D.H., 1981. Generalizing the Hough transform to detect arbitrary shapes. *Pattern Recognition* 13, 111–122.
- Barratt, D.C., Chan, C.S., Edwards, P.J., Penney, G.P., Slomczykowski, M., Carter, T.J., Hawkes, D.J., 2008. Instantiation and registration of statistical shape models of the femur and pelvis using 3D ultrasound imaging. *Medical Image Analysis* 12, 358–374.
- Bay, H., Tuytelaars, T., Van Gool, L., 2006. Surf: Speeded up robust features, in: *Computer vision—ECCV 2006*. Springer, pp. 404–417.
- Berger, M., Tagliasacchi, A., Seversky, L., Alliez, P., 2014. State of the art in surface reconstruction from point clouds. *EUROGRAPHICS STAR*.
- Bernard, F., Gemmar, P., Hertel, F., Goncalves, J., Thunberg, J., 2015a. Linear Shape Deformation Models with Local Support Using Graph-based Structured Matrix Factorisation. *arXiv:1510.08291*.
- Bernard, F., Gemmar, P., Husch, A., Hertel, F., 2012. Improvements on the feasibility of active shape model-based subthalamic nucleus segmentation. *Biomedical Engineering/Biomedizinische Technik* 57, 42–45.
- Bernard, F., Gemmar, P., Husch, A., Saleh, C., Neb, H., Dooms, G., Hertel, F., 2014. Improving the Consistency of Manual Deep Brain Structure Segmentations by Combining Variational Interpolation, Simultaneous Multi-Modality Visualisation and Histogram Equilisation. *Biomedical Engineering / Biomedizinische Technik* 59, 131–134.
- Bernard, F., Salamanca, L., Thunberg, J., Hertel, F., Goncalves, J., Gemmar, P., 2015b. Shape-aware 3D Interpolation using Statistical Shape Models, in: *Proceedings of Shape Symposium, DeleMont*.
- Bernard, F., Vlassis, N., Gemmar, P., Husch, A., Thunberg, J., Goncalves, J., Hertel, F., 2016. Fast correspondences for statistical shape models of brain structures, in: *Proc. SPIE Medical Imaging, San Diego*.
- Bernardini, F., Mittleman, J., Rushmeier, H., Silva, C., Taubin, G., 1999. The ball-pivoting algorithm for surface reconstruction. *IEEE Transactions on Visualization and Computer Graphics* 5, 349–359.
- Besl, P.J., McKay, N.D., 1992. A method for registration of 3-D shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 14, 239–256.
- Billings, S.D., Boctor, E.M., Taylor, R.H., 2015. Iterative most-likely point registration (IMLP): a robust algorithm for computing optimal shape alignment. *PloS one* 10.
- Blanc, R., Székely, G., 2012. Confidence regions for statistical model based shape prediction from sparse observations. *IEEE Transactions on Medical Imaging* 31, 1300–1310.
- Blanz, V., Mehl, A., Vetter, T., Seidel, H.P., 2004. A statistical method for robust 3D surface reconstruction from sparse data, in: *3D Data Processing, Visualization and Transmission*, pp. 293–300.
- Bolle, R.M., Vemuri, B.C., 1991. On three-dimensional surface reconstruction methods. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 13, 1–13.
- Bouaziz, S., Tagliasacchi, A., Pauly, M., 2013. Sparse iterative closest point. *Computer Graphics Forum* 32, 1–11.
- Boyd, S., Vandenberghe, L., 2009. *Convex Optimization*. Cambridge University Press.
- de Bruijne, M., van Ginneken, B., Viergever, M.A., Niessen, W.J., 2004. Interactive segmentation of abdominal aortic aneurysms in CTA images. *Medical Image Analysis* 8, 127–138.
- Chan, C.S., Edwards, P.J., Hawkes, D.J., 2003. Integration of ultrasound-based registration with statistical shape models for computer-assisted orthopaedic surgery, in: *SPIE Medical Imaging*, pp. 414–424.
- Chang, W., Zwicker, M., 2009. Range Scan Registration Using Reduced Deformable Models. *Computer Graphics Forum* 28, 447–456.
- Cootes, T.F., Taylor, C.J., 1992. Active Shape Models - Smart Snakes, in: *British Machine Vision Conference*, pp. 266–275.
- Dempster, A.P., Laird, N.M., Rubin, D.B., 1977. Maximum Likelihood from Incomplete Data via the EM Algorithm. *Journal of the Royal Statistical Society. Series B (Methodological)* 39, 1–38.
- Edelsbrunner, H., Mücke, E.P., 1994. Three-dimensional alpha shapes. *ACM Transactions on Graphics* 13, 43–72.
- Fleute, M., Lavallée, S., 1998. Building a complete surface model from sparse data using statistical shape models: Application to computer assisted knee surgery, in: *MICCAI*.
- Fleute, M., Lavallée, S., Julliard, R., 1999. Incorporating a statistically based shape model into a system for computer-assisted anterior cruciate ligament surgery. *Medical Image Analysis* 3, 209–222.
- Fonov, V.S., Evans, A.C., McKinstry, R.C., Almlí, C.R., Collins, D.L., 2009. Unbiased nonlinear average age-appropriate brain templates from birth to adulthood. *Neuroimage* 47, S102.
- Gal, R., Shamir, A., Hassner, T., Pauly, M., Cohen Or, D., 2007. Surface reconstruction using local shape priors., in: *Symposium on Geometry Processing*, pp. 253–262.
- van Ginneken, B., de Bruijne, M., Loog, M., Viergever, M.A., 2003. Interactive shape models, in: *SPIE Medical Imaging*, pp. 1206–1216.
- van Ginneken, B., Frangi, A.F., Frangi, R.F., Staal, J.J., Ter Haar Romeny, B.M., Viergever, M.A., 2002. Active Shape Model Segmentation with Optimal Features. *IEEE Transactions on Medical Imaging* 21, 924–933.
- Guo, G., Jiang, T., Wang, Y., Gao, W., 2012. Recovering Missing Contours for Occluded Object Detection. *IEEE Signal Processing Letters* 19, 463–466.
- Guo, G., Jiang, T., Wang, Y., Gao, W., 2013. 2-D shape completion with shape priors. *Chinese Science Bulletin* 58, 3430–3436.
- Heckel, F., Konrad, O., Karl Hahn, H., Peitgen, H.O., 2011. Interactive 3D medical image segmentation with energy-minimizing implicit functions. *Computers & Graphics* 35, 275–287.
- Heimann, T., Meinzer, H.P., 2009. Statistical shape models for 3D medical image segmentation: A review. *Medical Image Analysis* 13, 543–563.
- van den Hengel, A., Dick, A.R., Thormählen, T., Ward, B., Torr, P.H.S., 2007. Interactive 3D Model Completion, in: *Digital Image Computing Techniques and Applications*, pp. 175–181.
- Herman, G.T., Zheng, J., Bucholtz, C.A., 1992. Shape-based interpolation. *IEEE Computer Graphics and Applications* 12, 69–79.
- Hill, A., Cootes, T.F., Taylor, C.J., 1995. Active Shape Models and the Shape Approximation Problem. *Image and Vision Computing* 14, 601–607.
- Hoppe, H., DeRose, T., Duchamp, T., McDonald, J.A., Stuetzle, W., 1992. Surface reconstruction from unorganized points. *SIGGRAPH*, 71–78.
- Horaud, R., Forbes, F., Yguel, M., Dewaele, G., Zhang, J., 2011. Rigid and articulated point registration with expectation conditional maximization. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 33, 587–602.
- Jenkinson, M., Smith, S., 2001. A global optimisation method for robust affine registration of brain images. *Medical Image Analysis* 5, 143–156.
- Kazhdan, M., Bolitho, M., Hoppe, H., 2006. Poisson surface recon-

- struction, in: Eurographics Symposium on Geometry Processing. Liu, J., Udupa, J.K., 2009. Oriented active shape models. *IEEE Transactions on Medical Imaging* 28, 571–584.
- Lowe, D.G., 1999. Object recognition from local scale-invariant features, in: *International Conference on Computer Vision*, pp. 1150–1157.
- Lu, M., Zheng, B., Takamatsu, J., Nishino, K., 2011. 3D shape restoration via matrix recovery, in: *Computer Vision–ACCV Workshops*, pp. 306–315.
- Maier-Hein, L., Franz, A.M., dos Santos, T.R., Schmidt, M., Fangerau, M., Meinzer, H.P., Fitzpatrick, J.M., 2012. Convergent iterative closest-point algorithm to accomodate anisotropic and inhomogeneous localization error. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 34, 1520–1532.
- Meng, X.L., Rubin, D.B., 1993. Maximum likelihood estimation via the ECM algorithm: A general framework. *Biometrika* 80, 267–278.
- Myronenko, A., Song, X., 2010. Point Set Registration: Coherent Point Drift. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 32, 2262–2275.
- Myronenko, A., Song, X., Carreira-Perpinán, M.A., 2007. Non-rigid point set registration: Coherent Point Drift, in: *Neural Information Processing Systems*, pp. 1009–1016.
- Nocedal, J., Wright, S.J., 2006. *Numerical Optimization*. Springer.
- Osada, R., Funkhouser, T., Chazelle, B., Dobkin, D., 2002. Shape distributions. *ACM Transactions on Graphics (TOG)* 21, 807–832.
- Park, S., Guo, X., Shin, H., Qin, H., 2006. Surface completion for shape and appearance. *The Visual Computer* 22, 168–180.
- Pauly, M., Mitra, N.J., Giesen, J., Gross, M.H., Guibas, L.J., 2005. Example-Based 3D Scan Completion., in: *Symposium on Geometry Processing*, pp. 23–32.
- Rajamani, K.T., Styner, M.A., Talib, H., Zheng, G., Nolte, L.P., Ballester, M.A.G., 2007. Statistical deformable bone models for robust 3D surface extrapolation from sparse data. *Medical Image Analysis* 11, 99–109.
- Raya, S.P., Udupa, J.K., 1990. Shape-based interpolation of multi-dimensional objects. *IEEE Transactions on Medical Imaging* 9, 32–42.
- Rusinkiewicz, S., Levoy, M., 2001. Efficient variants of the ICP algorithm, in: *3-D Digital Imaging and Modeling*.
- Schroers, C., Setzer, S., Weickert, J., 2014. A Variational Taxonomy for Surface Reconstruction from Oriented Points, in: *Computer Graphics Forum*, pp. 195–204.
- Schwarz, T., Heimann, T., Tetzlaff, R., Rau, A.M., Wolf, I., Meinzer, H.P., 2008. Interactive surface correction for 3D shape based segmentation, in: *SPIE Medical Imaging*, pp. 69143O–69143O.
- Segal, A., Haehnel, D., Thrun, S., 2009. Generalized-icp, in: *Robotics: Science and Systems*.
- Seim, H., Kainmüller, D., Lamecker, H., Bindernagel, M., Malinowski, J., Zachow, S., 2010. Model-based Auto-Segmentation of Knee Bones and Cartilage in MRI Data. *Proc. of Medical Image Analysis for the Clinic: A Grand Challenge*, 215–223.
- Shen, C.H., Fu, H., Chen, K., Hu, S.M., 2012. Structure recovery by part assembly, in: *SIGGRAPH*.
- Stoll, C., Karni, Z., Rössl, C., Yamauchi, H., Seidel, H.P., 2006. Template deformation for point cloud fitting. *SPBG*, 27–35.
- Tan, J.H., Acharya, U.R., 2014. Active spline model: A shape based model—interactive segmentation. *Digital Signal Processing* 35, 64–74.
- Timinger, H., Pekar, V., von Berg, J., Dietmayer, K., Kaus, M., 2003. Integration of interactive corrections to model-based segmentation algorithms, in: *Bildverarbeitung für die Medizin 2003*. Springer, pp. 171–175.
- Treece, G.M., Prager, R.W., Gee, A.H., 2000. Surface interpolation from sparse cross sections using region correspondence. *IEEE Transactions on Medical Imaging* 11, 1106–1114.
- Turk, G., O’Brien, J.F., 1999. Shape transformation using variational implicit functions, in: *SIGGRAPH*, pp. 335–342.
- Xue, N., Doellinger, M., Ho, C.P., Surowiec, R.K., Schwarz, R., 2015. Automatic detection of anatomical landmarks on the knee joint using mri data. *Journal of Magnetic Resonance Imaging* 41, 183–192.
- Yang, D., Zhang, S., Yan, Z., Tan, C., Li, K., Metaxas, D., 2015. Automated anatomical landmark detection on distal femur surface using convolutional neural network, in: *International Symposium on Biomedical Imaging (ISBI)*, pp. 17–21.
- Zheng, G., Rajamani, K.T., Nolte, L.P., 2006. Use of a Dense Surface Point Distribution Model in a Three-Stage Anatomical Shape Reconstruction from Sparse Information for Computer Assisted Orthopaedic Surgery: A Preliminary Study. *Computer Vision–ACCV*.

Appendix A. Partial Contours \mathcal{P}

```

Input:
 $I : \mathbb{R}^3 \supset \Omega \rightarrow \hat{\mathcal{L}}$  (multi-label segmentation image),
 $n_{\text{obj}}, n_{\text{cont}},$ 
 $n_{\mathcal{L}}, \mathcal{L}' \subseteq \hat{\mathcal{L}}$ 
 $\bar{\sigma}$ 
Output:  $\mathcal{P}$ 
Initialise:  $\mathcal{P} = \emptyset$ 
// select  $n_{\mathcal{L}}$  elements from  $\mathcal{L}'$ 
1  $\mathcal{L} = \text{randomSelect}(\mathcal{L}', n_{\mathcal{L}})$ 
// fill  $\mathcal{L}$  such that  $|\mathcal{L}| = n_{\text{obj}}$ 
2  $\mathcal{L} = \mathcal{L} \cup \text{randomSelect}(\hat{\mathcal{L}} \setminus \mathcal{L}', n_{\text{obj}} - n_{\mathcal{L}})$ 
// for all objects in  $\mathcal{L}$ 
3 for  $\ell \in \mathcal{L}$  do
    // get binary image for object  $\ell$ 
4  $I_{\ell} = (I == \ell)$ 
    // get point-cloud of object voxels
5  $\mathbf{Z} = \text{getPointCloud}(I_{\ell})$  //  $\mathbf{Z} \in \mathbb{R}^3 \times 3$ 
6  $c = \text{mean}(\mathbf{Z}, 1)$  // get centre-of-gravity
    // create  $n_{\text{cont}}$  contours on different slice
    // orientations
7  $\mathcal{O} = \emptyset$  // save previous orientations
8 for  $i \in \{1, \dots, n_{\text{cont}}\}$  do
    // select slice orientation
9 if  $\mathcal{O} = \{1, 2, 3\}$  then
    |  $\mathcal{O} = \emptyset$  // reset  $\mathcal{O}$ 
10 |  $o = \text{randomSelect}(\{1, 2, 3\} \setminus \mathcal{O}, 1)$ 
11 |  $\mathcal{O} = \mathcal{O} \cup o$ 
    // extract boundary of 2d slice
12 | do
    | // randomly select a slice index
13 |  $s = \text{round}(c(o) + \mathcal{N}(0, \text{std}(\mathbf{Z}_{:,o})/2))$ 
14 |  $J_{\ell} = \text{extract2dSlice}(I_{\ell}, s, o)$ 
15 |  $\mathcal{B} = \text{extract2dBoundary}(J_{\ell})$ 
16 | while  $\mathcal{B}$  is invalid boundary
17 | // add noise to  $\mathcal{B}$ 
18 |  $\mathcal{B} = \text{translateBoundaryByRandomVector}(\mathcal{B}, \bar{\sigma})$ 
    // extract random subcontour of  $\mathcal{B}$ 
19 |  $\mathcal{D} = \text{extract2dSubcontour}(\mathcal{B}, 0.5)$  //  $|\mathcal{C}| = 0.5|\mathcal{B}|$ 
    // convert contour to points
20 |  $\mathcal{D} = \text{subsample}(\mathcal{D})$ 
    // convert 2d contour to 3d
21 |  $\mathcal{C} = \text{projectFrom2dBackTo3d}(\mathcal{D}, s, o)$ 
22 |  $\mathcal{P} = \mathcal{P} \cup \mathcal{C}$ 

```

Algorithm 4: Pseudocode for the generation of partial contour samples of \mathcal{P} for the brain structure dataset. The main idea is to randomly select a 2D slice of the binary 3D segmentation image for a particular object. From the 2D slice of the ground truth segmentation a subcontour of the entire boundary is randomly selected. Eventually, combining the chosen slice index with the subcontour leads to a planar 3D contour that is then subsampled and the points are added to \mathcal{P} . “randomSelect(\mathcal{X}, n)” performs the random selection of n objects from the set \mathcal{X} .

Appendix B. Detailed Brain Shapes Results

In addition to the methods ICP, ISO and ANISO, we evaluate the ISO and ANISO methods using a downsampled PDM, which we refer to as ISO-ds and ANISO-ds. The purpose of the downsampled PDM is reduce the processing time. As such, the objective of the experiments with the methods that use a downsampled PDM is to evaluate the reconstruction quality that can be obtained when a fast processing is essential, such as for example required for interactive segmentation. To downsample the PDM we used a subset of the original PDM points with $\approx 0.2N$ vertices. Such a downsampled PDM is shown in Fig. 5 (b).

Each subplot in Figs. B.11-B.20 shows the processing time on the horizontal axis versus one of the five metrics described in Section 7.1 on the vertical axis. The five colours indicate the five evaluated methods ICP, ISO, ANISO, ISO-ds, ANISO-ds, where a small dot indicates the outcome of a single run and the big dot indicates the mean of all runs for a particular method. In each subplot for each of the five colours there are $17 \cdot 20 = 340$ small dots. Thus, each of the subplots depicts a total $5 \cdot 340 = 1700$ runs of the fitting methods. The total number of surface reconstructions that have been evaluated is $2 \cdot 4 \cdot 6 \cdot 1700 = 81.600$ (LOO and LAI, 4 degrees of noise, 6 settings of \mathcal{P} , where 4 settings of random points on the shape surface and 2 settings of partial contours were considered).

Each row in Figs. B.11-B.20 shows a different setting of \mathcal{P} , where the first four rows show various numbers P of randomly drawn points on the shape surface (cf. Section 7), and the last two rows show the two settings of the sparse points based on partial contours, as described in Section 7.2. Each column in Figs. B.11-B.20 shows a different degree of noise in \mathcal{P} (cf. Fig. 5 (a,b)).

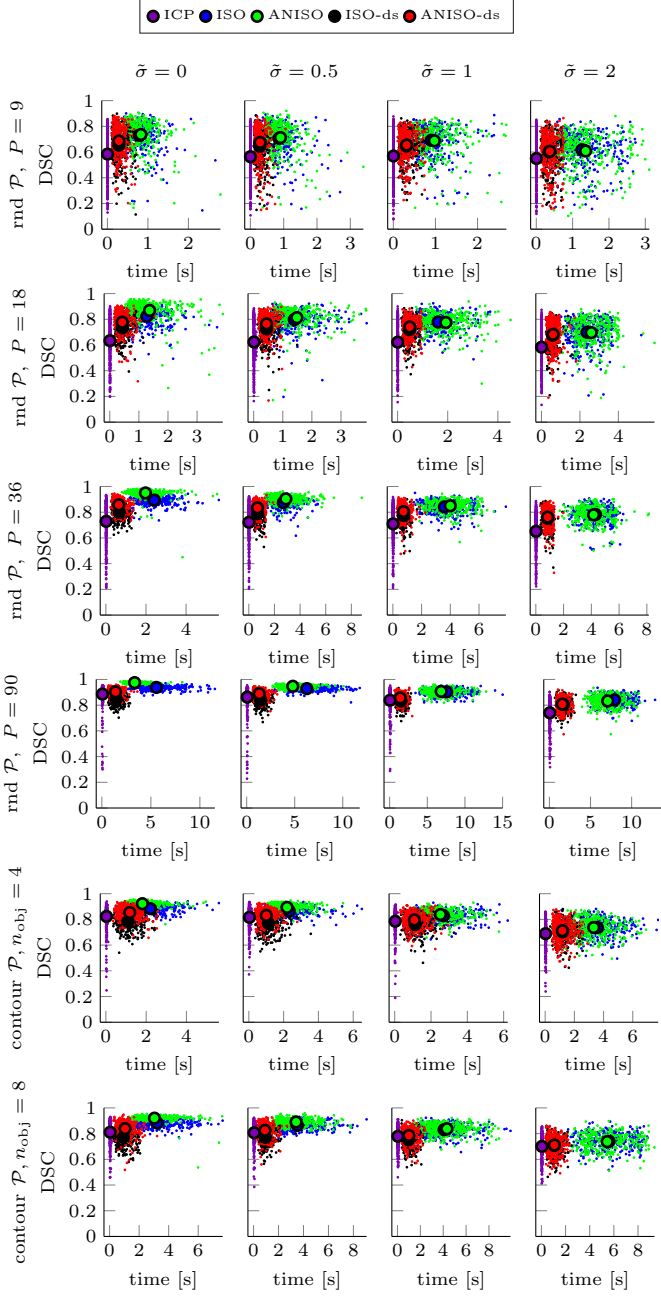


Figure B.11: DSC for brain shape LAI experiments.

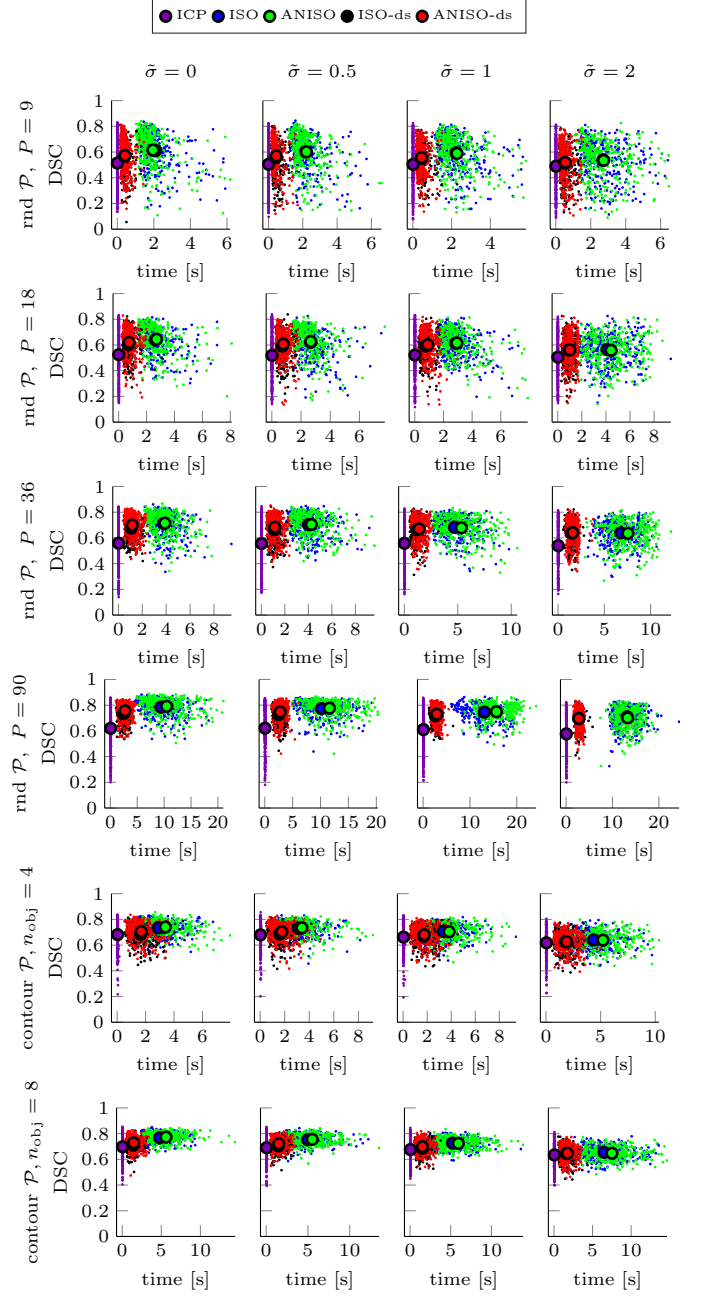


Figure B.12: DSC for brain shape LOO experiments.

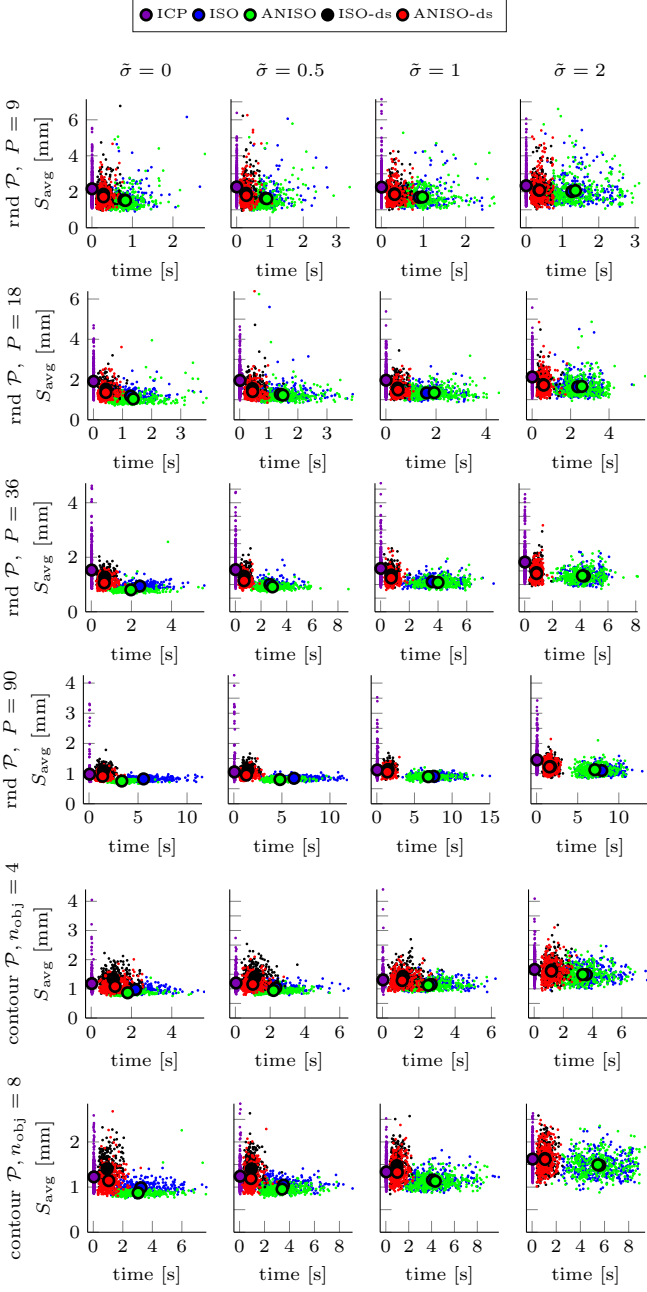


Figure B.13: S_{avg} for brain shape LAI experiments.

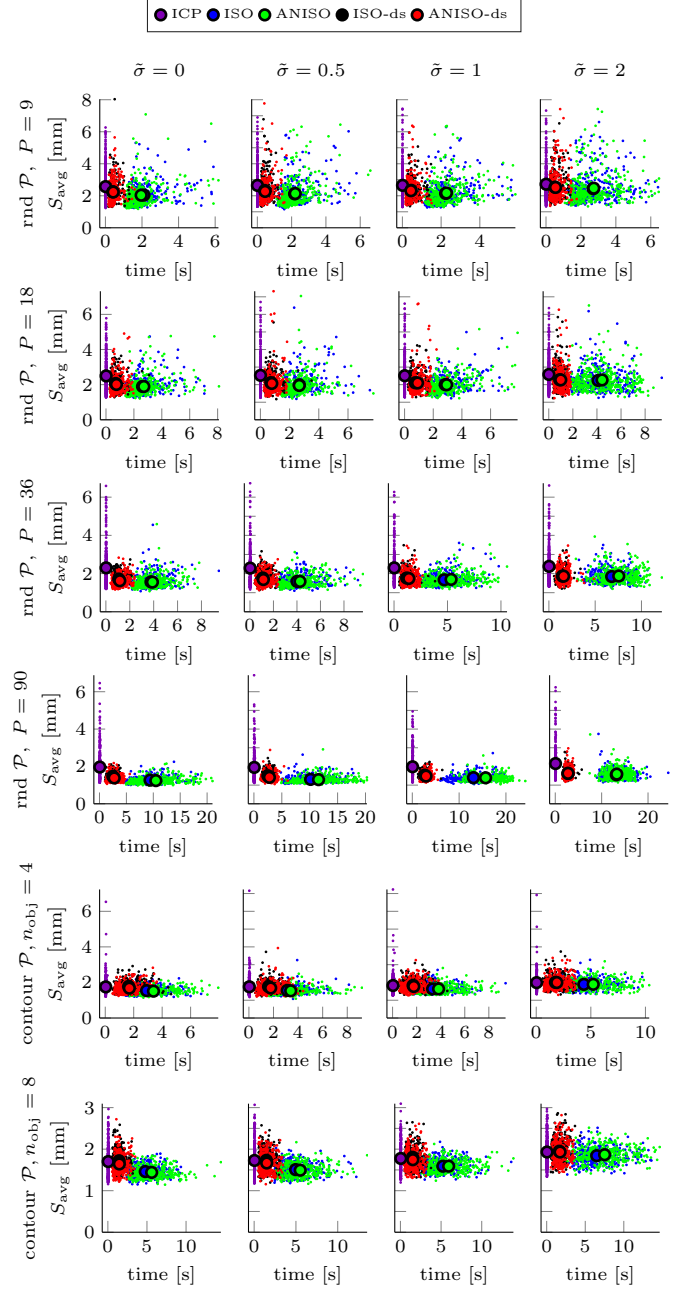


Figure B.14: S_{avg} for brain shape LOO experiments.

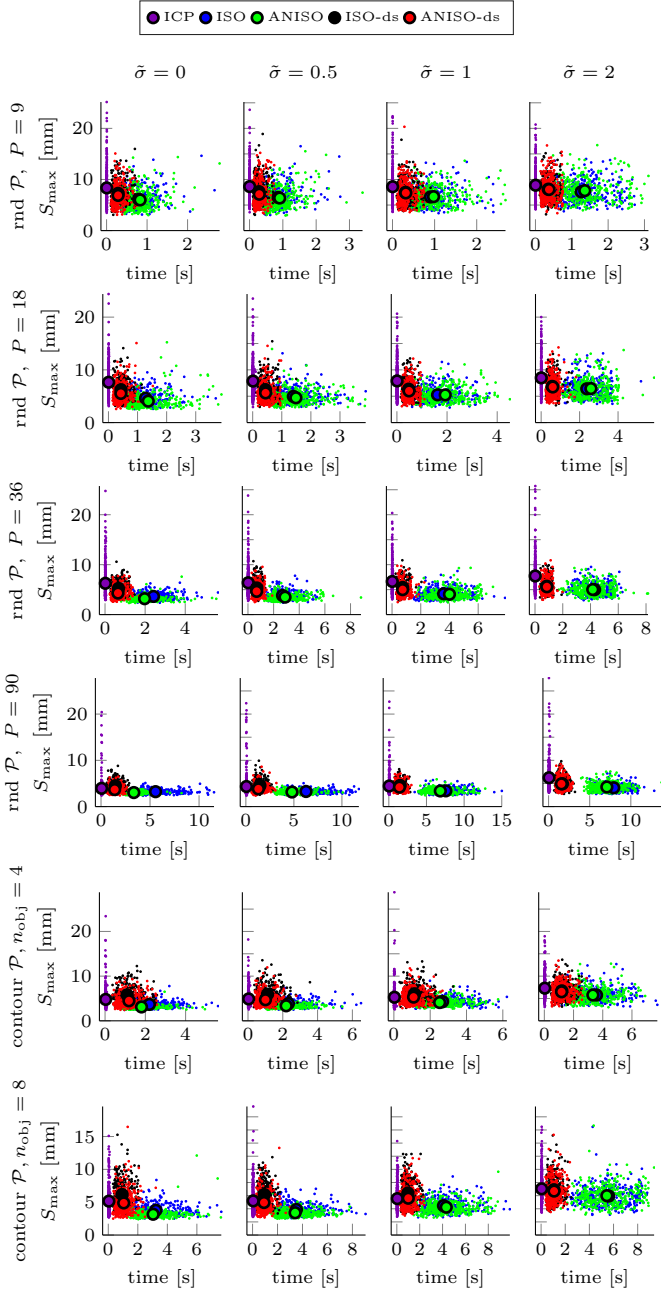


Figure B.15: S_{\max} for brain shape LAI experiments.

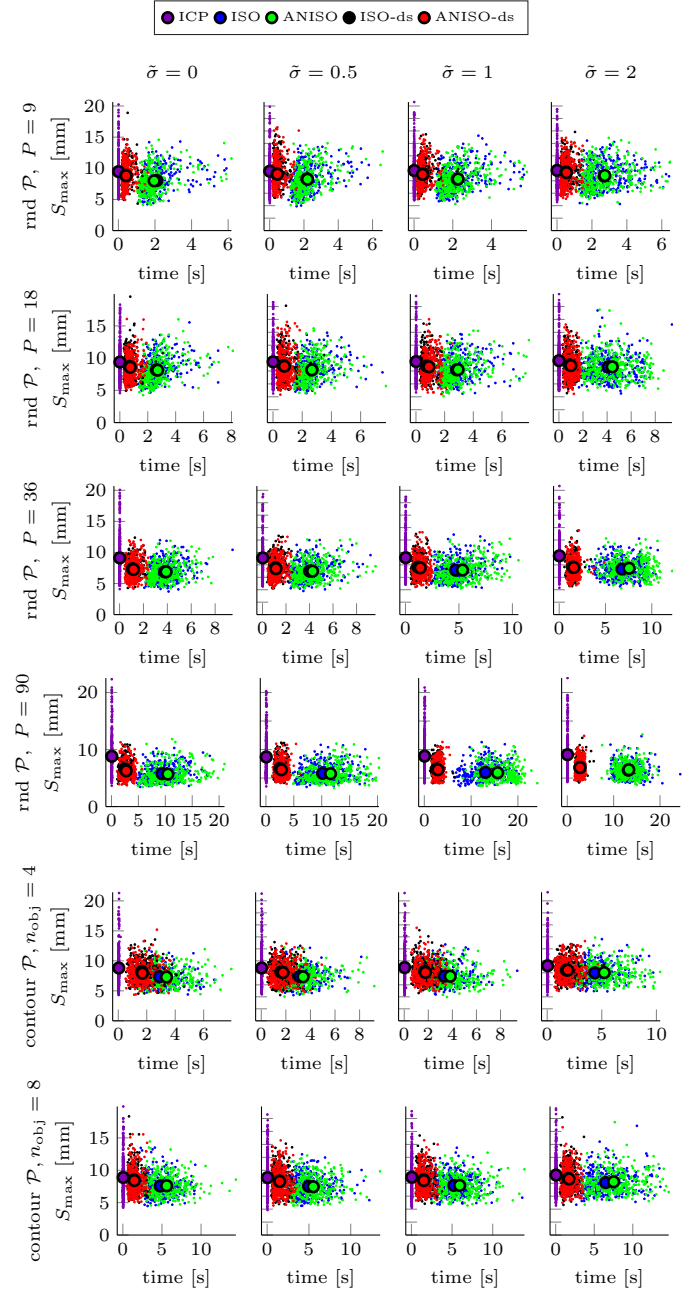


Figure B.16: S_{\max} for brain shape LOO experiments.




Figure B.17: D_{avg} for brain shape LAI experiments.

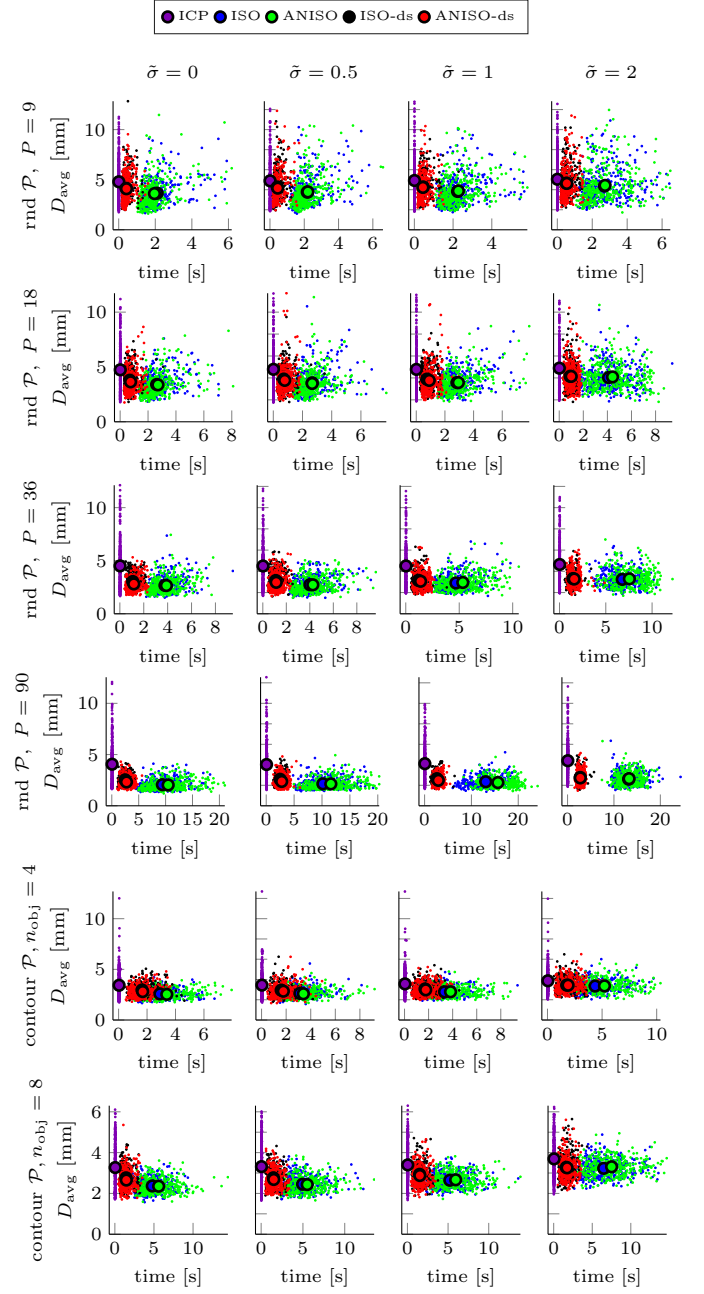


Figure B.18: D_{avg} for brain shape LOO experiments.




Figure B.19: D_{\max} for brain shape LAI experiments.

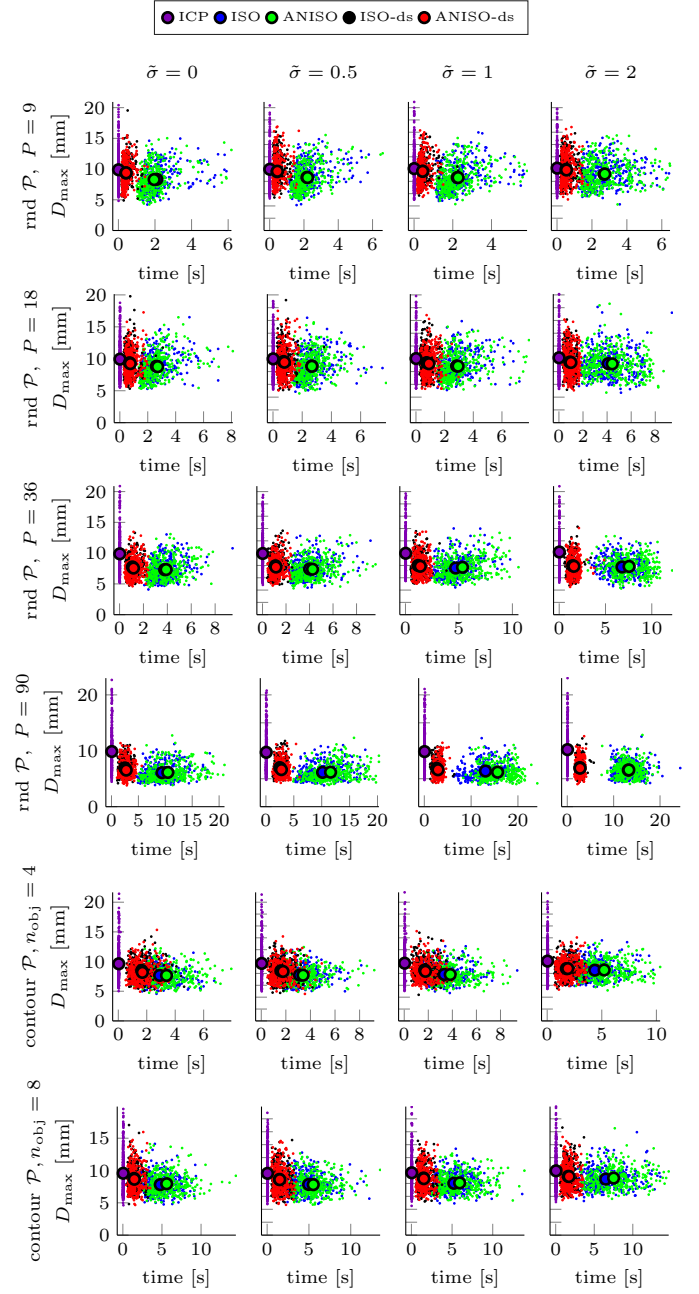


Figure B.20: D_{\max} for brain shape LOO experiments.

Appendix C. Detailed Femur Results

Each subplot in Figs. C.21, C.22 shows the processing time on the horizontal axis versus one of the five metrics described in Section 7.1 on the vertical axis. The five colours indicate the five evaluated methods ICP, ISO, ANISO, ISO-ds, ANISO-ds, where a small dot indicates the outcome of a single run and the big dot indicates the mean of all runs for a particular method. In each subplot for each of the five colours there are $60 \cdot 10 = 600$ small dots ($K = 60$ training shapes, 10 samples of \mathcal{P}). Thus, each of the subplots depicts a total $5 \cdot 600 = 3000$ runs of the fitting methods. The total number of surface reconstructions that are shown in Figs. C.21, C.22 is $2 \cdot 4 \cdot 3000 = 24000$ (LOO and LAI, 4 choices of P).

Each row in Figs. C.21, C.22 shows a different setting of \mathcal{P} . Each column in Figs. C.21, C.22 shows a different metric (cf. Section 7.1).

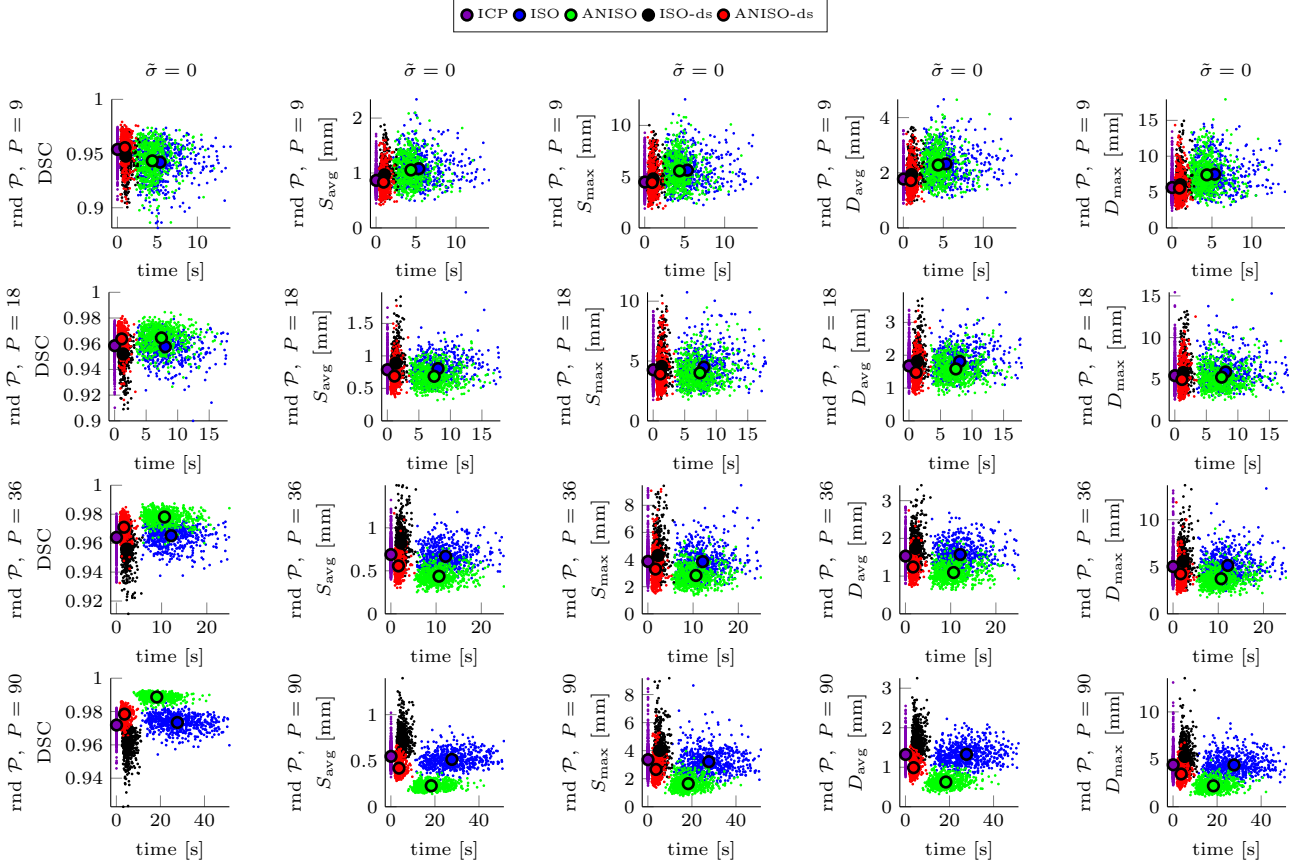


Figure C.21: Results for femur LAI experiments.

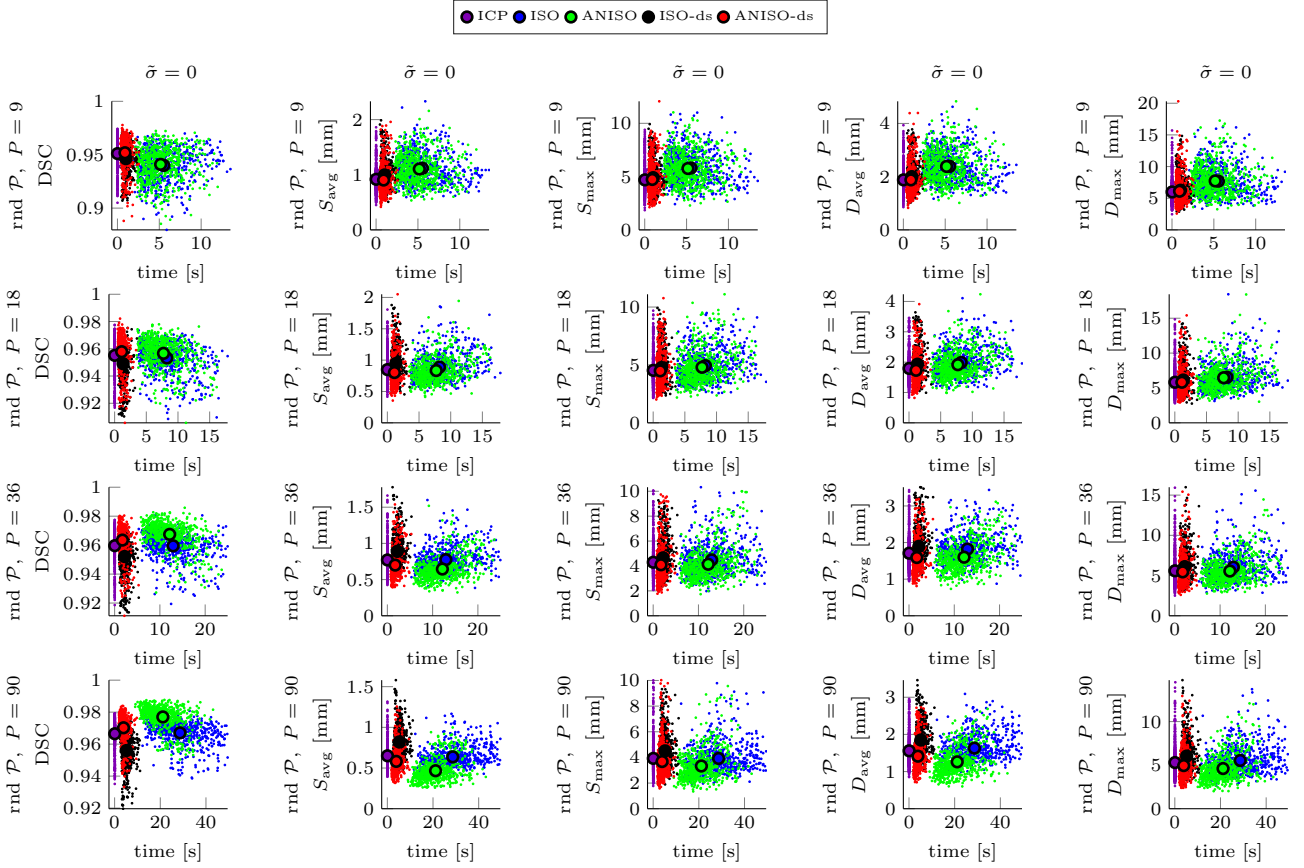


Figure C.22: Results for femur LOO experiments.